



DataDirect®

SequeLink®

Getting Started

© 2000 MERANT. All rights reserved. Printed in the U.S.A.

DataDirect, INTERSOLV, MicroFocus, Middleware, Net Express, PVCS, SequeLink, and TechGnosis are registered trademarks, and Client/Server MiddleWare, DataDirect Connect ADO, DataDirect Connect Integrator, DataDirect Connect JDBC, DataDirect Connect ODBC, DataDirect Connect OLE DB, DataDirect Connect Premium, DataDirect Reflector, DataDirect SequeLink Integrator, MERANT, PVCS Dimensions, MERANT, PVCS Metrics, PVCS Replicator, PVCS TeamLink, PVCS Tracker, PVCS TrackerLink, PVCS Version Manager, PVCS VM Server, and WebDBLink are trademarks of MERANT. All other trademarks are the property of their respective owners.

ACKNOWLEDGEMENT. PVCS® Dimensions™ is implemented using the ORACLE® Relational database management system. ORACLE is a registered trademark of Oracle Corporation, Redwood City, California.

Portions created by Netscape are Copyright (C) 1998-1999 Netscape Communications Corporation. All Rights Reserved.

Portions created by Eric Young are Copyright (C) 1995-1998 Eric Young. All Rights Reserved.

No part of this publication, with the exception of the software product user documentation contained on a CD-ROM, may be copied, photocopied, reproduced, transmitted, transcribed, or reduced to any electronic medium or machine-readable form without prior written consent of MERANT.

Licensees may duplicate the software product user documentation contained on a CD-ROM, but only to the extent necessary to support the users authorized access to the software under the license agreement. Any reproduction of the documentation, regardless of whether the documentation is reproduced in whole or in part, must be accompanied by this copyright statement in its entirety, without modification.

U.S. GOVERNMENT RESTRICTED RIGHTS. It is acknowledged that the Software and the Documentation were developed at private expense, that no part is in the public domain, and that the Software and Documentation are Commercial Computer Software provided with RESTRICTED RIGHTS under Federal Acquisition Regulations and agency supplements to them. Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of The Rights in Technical Data and Computer Software clause at DFAR 252.227-7013 et. seq. or subparagraphs (c)(1) and (2) of the Commercial Computer Software Restricted Rights at FAR 52.227-19, as applicable. Contractor is MERANT, 9420 Key West Avenue, Rockville, Maryland 20850. Rights are reserved under copyright laws of the United States with respect to unpublished portions of the Software.

MERANT
9420 Key West Avenue
Rockville, Maryland 20850

Table of Contents

Preface	7
What Is DataDirect SequeLink?	7
Using This Book	7
SequeLink Documentation	8
Conventions Used in This Book.	9
Typographical Conventions.	9
Mouse Conventions	10
Keyboard Conventions	11
Environment-Specific Information	11
Ordering Printed Books.	12
Contacting Technical Support.	14
1 Introduction	17
New Features in This Release	19
Where to Start	20
About the SequeLink Components	21
SequeLink Client	21
SequeLink Server	22
SequeLink Manager.	22
Understanding the SequeLink Client	23
Understanding SequeLink Services.	24
Understanding the SequeLink Manager Tools.	25
SequeLink Environment	27
Providing ODBC, ADO, and JDBC Data Access.	28
Using SequeLink to Link an Application to a Data Store.	29

SequeLink Architecture	31
SequeLink Two-Tier Architecture.	31
SequeLink <i>n</i> -Tier Architecture	33
SequeLink Packages	34
2 Sample Scenarios.	35
Scenario 1	35
Scenario 2	38
3 Planning Your SequeLink Configuration	41
Information You Need Before You Configure	42
Summary of What You Must Configure	42
Planning Your Client Data Sources	43
Configuring a Connection to a SequeLink Server	44
Referencing a Server Data Source	45
Planning Your SequeLink ODBC Client and SequeLink ADO Client Configurations	46
Configuring Quick Install Images	46
Including Client Data Source Configurations in Quick Install Images	47
Planning SequeLink Data Access Services	47
Default Behavior on Windows NT, Windows 2000, and UNIX.	48
Default Behavior on OS/390	53
Examples of Changing Default Behavior	56
Planning Your Connection Model	58
Planning Security	59
About SequeLink Security	59
Planning Security on Windows and UNIX.	67
Planning Security on OS/390	68

Planning System Administration 70

 Local System Administration..... 70

 Remote System Administration 71

Planning Monitoring and Event Tracing 72

 Monitoring..... 72

 Event Tracing 73

Examples: Configuring SequeLink 74

 Creating a Server Data Source 77

 Modifying the Default Server Data Source 83

 Configuring ODBC Client Data Sources 84

4 Migrating to SequeLink 5.x 89

 4.5 Data Sources Vs. 5.x Data Sources 89

Glossary 91

Index 95

6 Table of Contents

Preface

This book is your guide to getting started with MERANT™ DataDirect® SequeLink® 5.1. Read on to find out more about your SequeLink environment and how to use this book.

What Is DataDirect SequeLink?

DataDirect SequeLink is a middleware product that provides point-to-point connections from client to server for the latest data access standards, including Open Database Connectivity (ODBC), Java Database Connectivity (JDBC) applications, and ActiveX Data Objects (ADO) applications.

Using This Book

This book assumes that you are familiar with your operating system and its commands; the concept of directories; the management of user accounts and security access; and your network protocol and its configuration.

This book contains the following information:

- [Chapter 1 “Introduction” on page 17](#) introduces some concepts that will help you use SequeLink to provide data access across your enterprise.
- [Chapter 2 “Sample Scenarios” on page 35](#) provides sample scenarios that describe how SequeLink might be used to implement data access for a data consumer application.

- [Chapter 3 “Planning Your SequeLink Configuration” on page 41](#) provides information you need to know as you plan your SequeLink configuration.
- [Chapter 4 “Migrating to SequeLink 5.x” on page 89](#) provides high-level information about migrating from SequeLink 4.5 to SequeLink 5.x.

NOTE: This book refers the reader to Web URLs for more information about specific topics, including Web URLs not maintained by MERANT. Because it is the nature of Web content to change frequently, MERANT can guarantee only that the URLs referenced in this book were correct at the time of publishing.

SequeLink Documentation

The following documentation is provided on your SequeLink CD in PDF format. You can view the online documentation on the CD using the Acrobat Reader.

The following table provides a guide for finding information in your SequeLink documentation.

For information about...	Go to...
SequeLink concepts and planning your SequeLink environment	<i>Getting Started with SequeLink</i>
Installing the SequeLink middleware components	<i>SequeLink Installation Guide</i>
Administering your SequeLink environment	<i>SequeLink Administrator's Guide</i>

For information about...	Go to...
Developing ODBC, ADO, and JDBC applications for the SequeLink environment	<i>SequeLink Developer's Reference</i>
Troubleshooting and referencing error messages	<i>SequeLink Troubleshooting Guide and Reference</i>

DataDirect product documentation is also available in PDF and HTML formats on the MERANT DataDirect Web site:

<http://www.merant.com/products/datadirect/download/docs/dochome.asp>

Conventions Used in This Book

This section describes the typography, terminology, and other conventions used in this book.

Typographical Conventions

This book uses the following typographical conventions:

Convention	Explanation
<i>italics</i>	Introduces new terms that you may not be familiar with, and is used occasionally for emphasis.
bold	Emphasizes important information. Also indicates button, menu, and icon names on which you can act. For example, click Next .

Convention	Explanation
UPPERCASE	Indicates the name of a file. For operating environments that use case-sensitive file names, the correct capitalization is used in information specific to those environments. Also indicates keys or key combinations that you can use. For example, press the ENTER key.
<code>monospace</code>	Indicates syntax examples, values that you specify, or results that you receive.
<i>monospaced italics</i>	Indicates names that are placeholders for values you specify; for example, <i>filename</i> .
forward slash /	Separates menus and their associated commands. For example, Select File / Copy means to select Copy from the File menu.
vertical rule	Indicates an OR separator to delineate items.
brackets []	Indicates optional items. For example, in the following statement: SELECT [DISTINCT], DISTINCT is an optional keyword.
braces { }	Indicates that you must select one item. For example, {yes no} means you must specify either yes or no.
ellipsis . . .	Indicates that the immediately preceding item can be repeated any number of times in succession. An ellipsis following a closing bracket indicates that all information in that unit can be repeated.

Mouse Conventions

This action...	Means to...
Click	Point to an object with the mouse pointer and press the left mouse button.
Double-click	Click the left mouse button twice.

This action...	Means to...
Right-click	Click the right mouse button.
Drag	Press and hold the left mouse button while dragging item(s) to another part of the screen.
SHIFT+Click	Press and hold the SHIFT key; then click a selection. This lets you select a series of objects.
CTRL+Click	Press and hold the CTRL key; then click a selection. This lets you select objects randomly.

Keyboard Conventions

Select menu items by using the mouse or pressing ALT+ the key letter of the menu name or item.

Environment-Specific Information

This book supports users of various operating environments. Where it provides information that does not apply to all supported environments, the following symbols are used to identify that information.

Symbol Environment



Windows. Information specific to the Microsoft Windows 95, Windows 98, Windows ME, Windows NT, and Windows 2000 environments is identified by the Windows symbol.



Windows NT. Information specific to the Microsoft Windows NT environment is identified by the Windows symbol and the letters "NT."

Symbol	Environment
	Windows 2000. Information specific to the Microsoft Windows 2000 environment is identified by the Windows symbol and the number "2000".
	<i>UNIX</i> . Information specific to UNIX environments is identified by this symbol, which applies to all supported UNIX environments. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Ltd.
OS/390	<i>OS/390</i> . Information specific to OS/390 environments is identified by the letters and numbers "OS/390".

Ordering Printed Books

As part of your SequeLink license agreement, you may print and distribute as many copies of the SequeLink books as needed.

If you do not want to print each of these online books, you can order printed versions from MERANT. To order, please complete the following order form and fax your request to MERANT at (919) 461-4526.

Contacting Technical Support

MERANT provides technical support for all registered users of SequeLink, including limited installation support, for the first 30 days. For support after that time, contact us using one of the following methods or purchase further support by enrolling in the SupportNet program. For more information about SupportNet, contact your sales representative.

The MERANT Web site provides the latest support information through SupportNet Online, our global service network that provides access to valuable tools and information. Our SupportNet users access information using the Web, automatic email notification, newsgroups, and regional user groups. SupportNet Online includes a knowledge base that allows you to search on keywords for technical bulletins and other information. You also can download product fixes for your DataDirect products.

World Wide Web

<http://support.merant.com>

E-Mail

USA, Canada, and Mexico	datadirect.answerline@merant.com
Australia and New Zealand	australia.answerline@merant.com
Japan	jpn.answerline@merant.co.jp
All other countries	int.datadirect.answerline@merant.com

Local Telephone Support

Local phone numbers can be found at:

<http://support.merant.com/websupport/contact/supportnetanswerline.asp>

Live Answerline telephone support is available 24 hours a day, seven days a week.

Fax and Mail Information

Fax US	1 919 461 4527
Fax International	+32 (0) 15 32 09 19
Mail	1500 Perimeter Park Drive, Suite 100, Morrisville, NC 27560 USA

When you contact us, please provide the following information:

- The **product serial number** located on the Product Registration Information card or on a product serial number card in your package. The number will be checked to verify your support eligibility. If you do not have a SupportNet contract, we will ask you to speak with a sales representative.
- Your **name and organization**. For a first-time call, you may be asked for full customer information, including location and contact details.
- The **version number** of your DataDirect product.
- The type and version of your **operating system**.
- Any **third-party software or other environment information** required to understand the problem.
- A **brief description of the problem**, including any error messages that you have received, **and the steps preceding the occurrence of the problem**. Depending on the complexity of the problem, you may be asked to submit an example so that we can recreate the problem.
- An assessment of the **severity level** of the problem.

1 Introduction

DataDirect SequeLink is a middleware product that provides point-to-point connections from client to server for the latest data access standards—Open Database Connectivity (ODBC 3.5), Java Database Connectivity (JDBC 2.0, including Optional Package), and ActiveX Data Objects (ADO 2.5). In addition, SequeLink allows you to centrally configure your data access environment and manage data access activity.

Today's complex Information Technology (IT) environments require data access components that provide superior interoperability, performance, and manageability. SequeLink middleware fulfills these requirements by providing the following key advantages for complex IT environments:

- **Data Connectivity.** SequeLink provides universal data connectivity for the latest ODBC, ADO, and JDBC standards to a variety of data stores, including mainframe data. Additionally, SequeLink supports distributed transactions for Microsoft Distributed Transaction Controller (MS DTC) and Java Transaction API (JTA). SequeLink's component implementation allows you to manage your entire data access environment regardless of the underlying operating systems on which the SequeLink components run. In addition, SequeLink Client is database independent—no extra client components are required if you decide to incorporate additional data store technologies in your data access infrastructure.
- **Interoperability.** SequeLink allows you to leverage existing and evolving technologies by adhering to industry standards rather than its own proprietary standards. In addition to supporting the latest data access standards, SequeLink allows you to use Lightweight Directory Access Protocol (LDAP) for centralized connection and configuration information and

Internet Inter-ORB Protocol (IIOP) for secure Internet and Intranet communication.

- **Security.** The messages between SequeLink middleware components that involve data requests and data transmitted over a network, Internet, or Intranet can be scrambled or, in some cases, encrypted. In addition, SequeLink supports authentication mechanisms provided by the database and/or by the operating system on which the SequeLink components run, such as NT Integrated Security on Windows and Resource Access Control Facility (RACF) on OS/390. SequeLink also supports read-only data store connections to keep the data in your data store secure from updates.
- **Systems management.** SequeLink delivers key Reliability, Availability, and Serviceability (RAS) by providing dynamic service attributes. The majority of SequeLink's service configuration attributes are dynamic, meaning that, if you change a setting for an attribute, the change takes affect immediately.
- **Scalability.** SequeLink provides superior performance and scalability through connection pooling at the client, through efficient use of server resources, and through its configurable multi-threaded implementation, which uses a dynamic worker thread model.

This chapter introduces some concepts that will help you use SequeLink to provide data access across your enterprise.

New Features in This Release

SequeLink provides the following standard features:

- Centralized management, administration, and monitoring capabilities using the SequeLink Manager
- Improved scalability and performance using a thread-pool engine within the SequeLink Server on Windows NT/2000 and UNIX
- Support for latest data access specifications, databases, and operating system versions

In addition, this SequeLink 5.1 release also provides the following new features:

- Support for DB2, Informix, and Sybase on Windows NT/2000 and UNIX. For a complete list of the databases supported by SequeLink 5.1, see:

<http://www.merant.com/products/datadirect/odbc/sl5/techdocs/matrix.asp>

- Enhancements to SequeLink Server on OS/390, including support for DB2 V6R1 and support for literals in stored procedure calls
- Improved support for distributed transactions (MS DTC and JTA)

Where to Start

The following table provides a guide to topics about the key features of SequeLink to help you plan, configure, and administer your SequeLink environment.

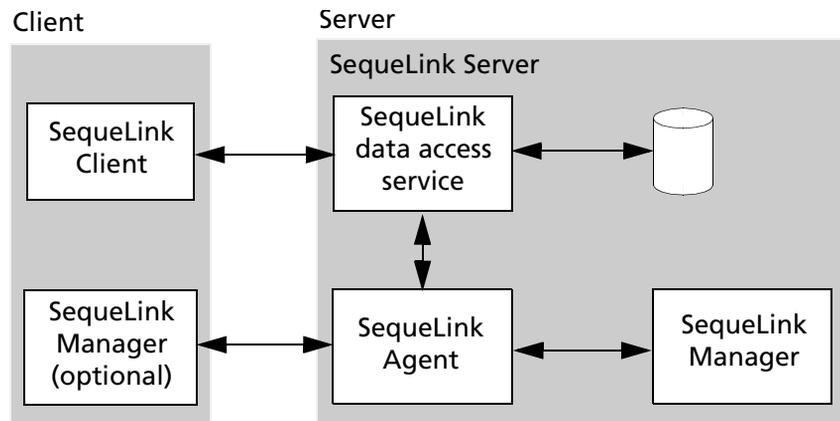
SequeLink Feature	See...
SequeLink Client	“Understanding the SequeLink Client” on page 23 to learn how SequeLink’s thin-client component can ease configuration and administration
SequeLink Server	“Understanding SequeLink Services” on page 24 to learn how SequeLink provides optimized data connectivity, performance, and administration
SequeLink Manager	“Understanding the SequeLink Manager Tools” on page 25 to learn about the type of tasks you can perform using the SequeLink Manager to configure and manage your SequeLink environment and monitor data access activity
SequeLink Server thread-pool engine on Windows NT/2000 and UNIX	“Planning Your Connection Model” on page 58 to learn how the thread-pool engine ensures efficient use of server resources
Security	“Planning Security” on page 59 for a discussion of the security options provided by SequeLink

About the SequeLink Components

SequeLink operates through the following middleware components shown in [Figure 1-1](#):

- SequeLink Client
- SequeLink Server
- SequeLink Manager

Figure 1-1. SequeLink Components



SequeLink Client

SequeLink Client uses a thin-client architecture to provide a single, universal interface for data access that is easy to install and requires “near-zero” administration. SequeLink has three different clients—the SequeLink ODBC Client supports ODBC applications, the SequeLink Java Client supports JDBC applications, and the SequeLink ADO Client supports ADO/OLE DB applications. The SequeLink Client is database

independent, so, if you decide to incorporate additional data store technologies, you do not need to update or install a new SequeLink Client.

For more information about the SequeLink Client, see [“Understanding the SequeLink Client” on page 23](#).

SequeLink Server

SequeLink Server installs the following server software service components to provide data connectivity, performance, and administration for two-tier client/server and n -tier web/application server environments:

- **SequeLink data access services** handle data access requests from any SequeLink Client. Multiple SequeLink data access services can run on the same SequeLink Server. For example, SequeLink Server for Oracle and SequeLink Server for Microsoft SQL Server can run side-by-side on the same machine.
- **SequeLink Agent services** carry out configuration, management, and monitoring requests from any SequeLink Manager. The SequeLink Agent can service multiple SequeLink services on the same SequeLink Server.

For more information about SequeLink services, see [“Understanding SequeLink Services” on page 24](#).

SequeLink Manager

SequeLink Manager is a tool that allows you to centrally configure, manage, and monitor your entire data access infrastructure. By default, this tool is installed on the server that contains the SequeLink Server software; optionally, you can install it on a networked client.

The SequeLink Manager is implemented differently depending on platform.

NOTE: Only SequeLink 5.1 services can be configured, managed, or monitored with the SequeLink Manager 5.1.

For more information about the SequeLink Manager, see [“Understanding the SequeLink Manager Tools” on page 25](#).

Understanding the SequeLink Client

SequeLink’s thin-client component drastically reduces the amount of configuration that you must initially perform and the amount of time you must spend to reconfigure your data access infrastructure when a server configuration changes. Client data sources contain minimal information; most configuration information resides on the server, resulting in "near-zero" client administration.

Also, SequeLink allows administrators to use Lightweight Directory Access Protocol (LDAP) directories for centralized information retrieval. A client data source can reference an LDAP directory to retrieve server connection information, which can reduce the time it takes to reconfigure your infrastructure when a change takes place. For example, if a database must be moved to a different server, the administrator does not have to reconfigure user applications or the client data sources that must now access the new server because the connection information is stored centrally in an LDAP directory. Therefore, the administrator only needs to update the LDAP directory entries to allow the SequeLink Clients to connect to the new server.

Understanding SequeLink Services

In this section, we take a closer look at SequeLink services and how they work. Remember that SequeLink services are components of SequeLink Server.

SequeLink data access services can handle data access requests from any SequeLink Client. A data access service “services” a specific type of data store (for example, Oracle). SequeLink provides an optimized data access service for each type of data store it supports. Multiple SequeLink data access services can run on the same server. When you complete the installation of the SequeLink Server software as documented in the *SequeLink Installation Guide*, a SequeLink data access service is configured for the type of SequeLink Server you installed (for example, SequeLink Server for Oracle).

When a SequeLink Client connects to a SequeLink data access service, the data access functionality of the session is governed by the attributes defined for the data access service. For example, if the service attribute `DataSourceReadOnly=Select`, the client application can perform only Select statements when using that service. See the *SequeLink Administrator’s Guide* for a complete list of service attributes.

SequeLink Agent services carry out configuration, management, and monitoring requests from any SequeLink Manager. When you complete the installation of the SequeLink Server software as documented in the *SequeLink Installation Guide*, a SequeLink Agent service is created, configured, and is ready for requests.

Understanding the SequeLink Manager Tools

The SequeLink Manager tool can be used to perform administrative and monitoring requests.

Administrative Requests

The type of administrative requests you can issue to a SequeLink Agent, and the SequeLink Manager tool you can use to issue the requests, depends on the platform you are administering—Windows, UNIX, or OS/390. The following list describes the types of administrative requests you can issue:

■ Configuration

- Creating and managing SequeLink services
- Creating and managing SequeLink server data sources
- Configuring monitoring profiles, which determines the data access events that can be monitored (viewed) using the SequeLink Manager
- Configuring event tracing profiles, which determines the data access events that are written to an event trace file

■ Management

- On Windows and UNIX only: Starting and stopping SequeLink services
- Stopping active data access user sessions
- Reviewing traced events to analyze a problem during an earlier data access activity

Monitoring Requests

The SequeLink Manager allows you to perform the following monitoring tasks:

- Viewing details about active services
- Viewing active user sessions and information about "live" data access activities

For example, you can easily view how many transactions have been processed or how many rows have been fetched by all user sessions. Also, if a user session is not performing correctly (such as the session keeps fetching thousands and thousands of rows), you can use the SequeLink Manager to identify and end that specific user session.

Additionally, the SequeLink Manager allows you to troubleshoot previous events. For example, if an error occurs during a nightly data processing job, you can look at an "event trace" to troubleshoot the problem.

SequeLink Manager Implementations

SequeLink provides the following implementations of the SequeLink Manager:

- **SequeLink Manager Snap-in** is a GUI designed as a snap-in to the Microsoft Management Console (MMC). It can be used to configure and manage SequeLink services on Windows and UNIX platforms. It can be used to monitor data access activity on Windows, UNIX, and OS/390 platforms.
- **SequeLink Manager Command-Line Tool** is a command-line interface that can be used to configure and manage SequeLink services on Windows and UNIX platforms. Similarly, it can be used to monitor data access activity on Windows and UNIX platforms.

- **SequeLink Manager for OS/390** is an ISPF dialog tool that can be used to configure and manage SequeLink services on the OS/390 platform, as well as monitor data access activity on the OS/390 platform. It can be installed only on OS/390 platforms.

[Table 1-1](#) shows the platforms on which you can install and run the different implementations of the SequeLink Manager.

Table 1-1. Platforms on which the SequeLink Manager Tools can be Installed

SequeLink Manager	Win NT/ 2000	UNIX	OS/390
SequeLink Manager Snap-in	✓		
SequeLink Manager Command-Line Tool	✓	✓	
SequeLink Manager for OS/390			✓

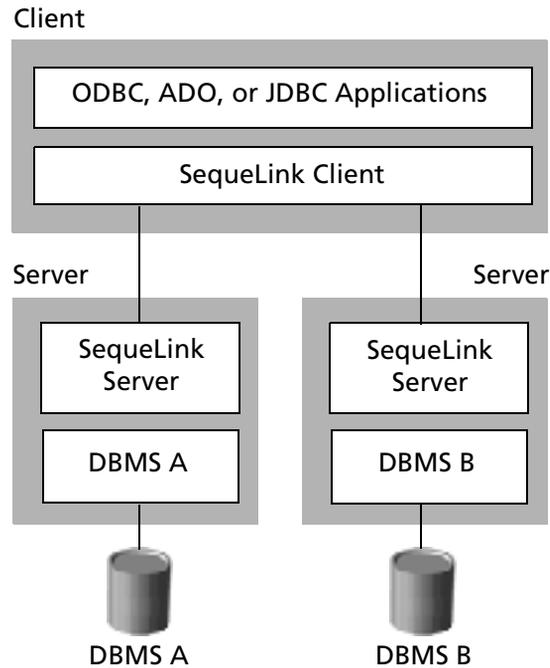
NOTE: Only SequeLink 5.1 services can be configured, managed, or monitored with the SequeLink Manager 5.1.

SequeLink Environment

Today's data access computing environments typically involve multiple and disparate data stores accessed over a variety of infrastructures. In addition, many businesses are relying more on the Internet to provide access to corporate data for their employees, customers, and partners.

SequeLink simplifies data access middleware requirements by providing data access for multiple data stores with a single client component as shown in [Figure 1-2 on page 28](#).

Figure 1-2. SequeLink Middleware Solution for Data Access

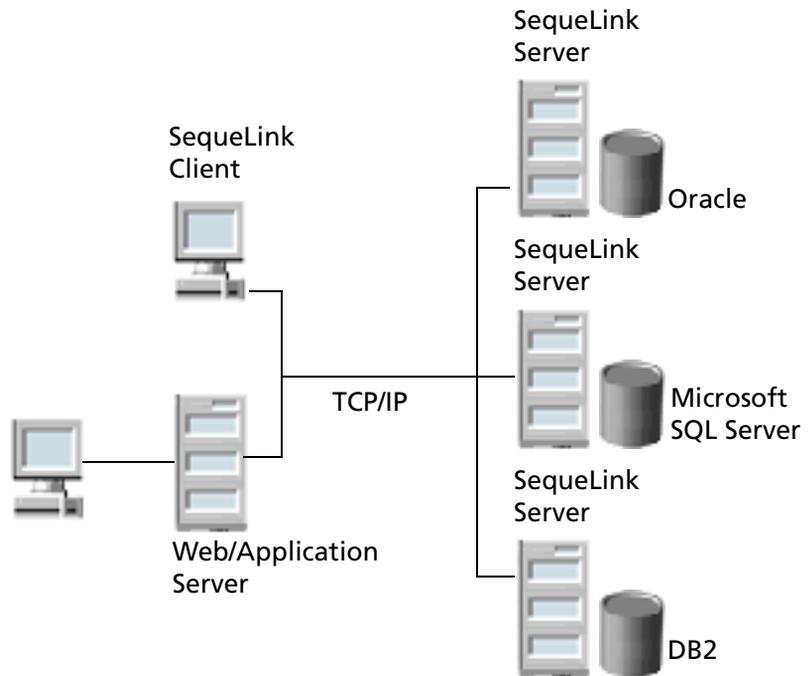


Providing ODBC, ADO, and JDBC Data Access

ODBC, ADO, and JDBC are application programming interfaces (APIs) that allow developers to develop, compile, and ship an application without targeting a specific type of data store. Developers can use SequeLink as the middleware that allows their applications to access data from a choice of supported data stores. SequeLink provides ODBC, ADO, and JDBC data access from a client workstation directly to a server running the

SequeLink Server software or from a web/application server to a server running SequeLink Server as shown in [Figure 1-3](#).

Figure 1-3. SequeLink Data Access in Client/Server and Web/Application Server Environments



Using SequeLink to Link an Application to a Data Store

Typically, a data access application accesses a SequeLink Server (specifically, a SequeLink data access service) by specifying a client data source that provides connection information to a specific server running SequeLink Server. SequeLink uses two

different types of data sources—server data sources and client data sources.

Server Data Sources

Server data sources are data sources configured on the SequeLink Server that contain settings that affect how the SequeLink service operates and settings that affect how data is accessed by SequeLink Clients. Centralizing this information on the server, instead of distributing it among hundreds of SequeLink Clients, provides easier management of your entire data access infrastructure. When you install SequeLink Server, a default server data source, named *Default*, is automatically created on the server. If necessary, you can modify the definition of the default server data source.

Client Data Sources

Client data sources are minimal data sources configured on the SequeLink Client that contain connection instructions to a SequeLink data access service. Client data sources are required when configuring the SequeLink ODBC Client or the SequeLink ADO Client. For SequeLink Java Clients, you can configure a client data source or a connection URL.

A client data source can be configured to retrieve connection information from an LDAP directory server. The connection information stored in an LDAP directory contains the IP address of the server that runs SequeLink Server; therefore, if the server IP address changes, you need only change the server in one place—the LDAP directory. You do not have to modify the SequeLink Client data sources or the user application.

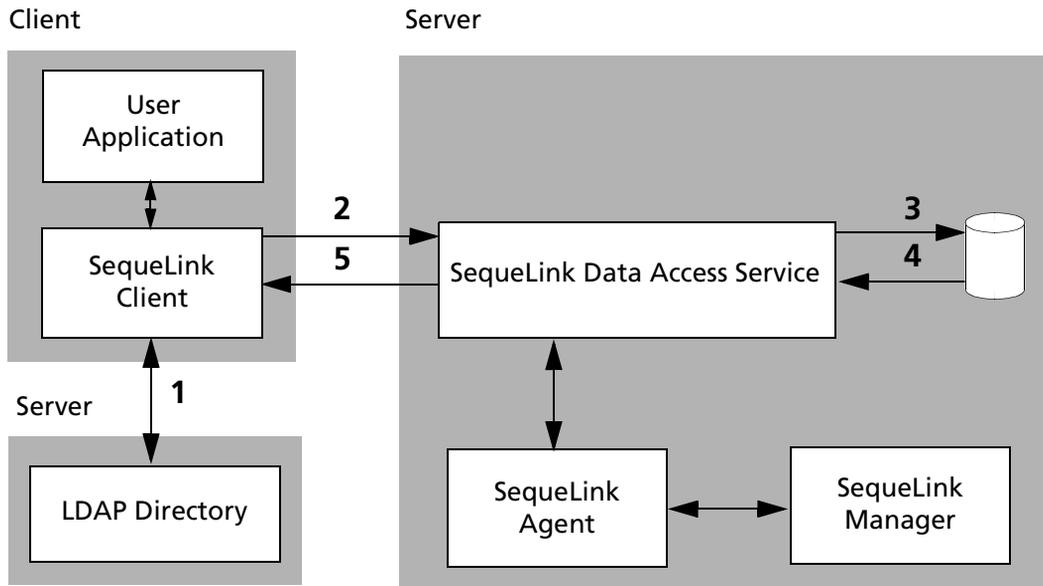
SequeLink Architecture

This section shows examples of SequeLink environments implemented with two-tier and n -tier architectures. These configurations explain the data access flow through the SequeLink components.

Additionally, the configurations show the SequeLink Manager and SequeLink Agent, which, together, allow the SequeLink administrator to control data access activities. For example, using the SequeLink Manager, an administrator can end an active data access user session. All actions the administrator performs on a data access service are handled by the SequeLink Agent.

SequeLink Two-Tier Architecture

SequeLink Server is often installed on the same server on which the database engine resides, and the SequeLink Client is often installed on the same workstation as the user application, as shown in [Figure 1-4 on page 32](#). This is a two-tier architecture because only two machines are needed for the configuration.

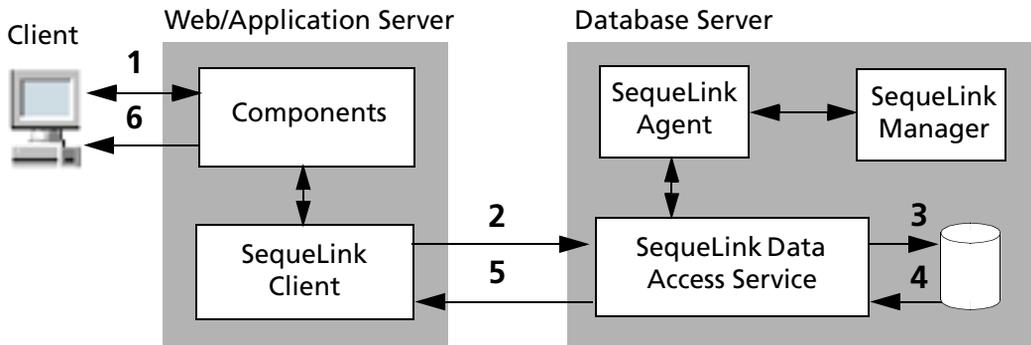
Figure 1-4. SequeLink Two-Tier Architecture**Data Access Architecture:**

- 1 SequeLink Clients can be configured to connect directly to a specific SequeLink Server or retrieve connection information from an LDAP directory.
- 2 User applications use the SequeLink Client to connect to the SequeLink Server. Applications make SQL calls from the SequeLink Client to the SequeLink Server using standard APIs, such as ODBC, ADO, or JDBC.
- 3 The SequeLink data access service passes the SQL request to the database engine.
- 4 The database engine processes the SQL request and passes results back to the SequeLink data access service.
- 5 The SequeLink data access service returns the results directly to the SequeLink Client and the user application.

SequeLink *n*-Tier Architecture

For maximum flexibility and centralized access, SequeLink Client can be installed on a middle-tier server between the client and a database server as shown in [Figure 1-5](#). An example of this configuration is when a client (tier 1) runs a web browser that downloads and displays a web page stored on the Web/Application Server (tier 2). On the web page, the user clicks a button that launches an application (component) on the Web/Application Server. This application uses ODBC, ADO, or JDBC to access a SequeLink Client that is also on the Web/Application Server. The SequeLink Client accesses data from the data store that is serviced by the SequeLink data access service residing on the Database Server (tier 3).

Figure 1-5. SequeLink *n*-Tier Architecture



Data Access Architecture:

- 1 In a three-tier architecture, a user application (for example, a web browser) in tier 1 may invoke components residing on the Web/Application Server that need to load the SequeLink Client to gain access to the data store on the Database Server.
- 2 SequeLink ODBC Client, SequeLink ADO Client, or SequeLink Java Client is installed on the Web/Application Server.

Applications make SQL calls from the Web/Application Server running the SequeLink Client to the SequeLink Server using standard APIs, such as ODBC, ADO, or JDBC.

- 3 The SequeLink data access service passes the SQL request to the database engine.
- 4 The database engine processes the SQL request and passes results back to the SequeLink data access service.
- 5 The SequeLink data access service returns the results to the Web/Application Server.
- 6 The client receives the results from the Web/Application Server.

SequeLink Packages

For the latest information about the operating system platforms, database management systems, and data access APIs supported by SequeLink, see:

<http://www.merant.com/products/datadirect/odbc/s15/techdocs/matrix.asp>

2 Sample Scenarios

This chapter provides sample scenarios that describe how SequeLink might be used to implement data access for a data consumer application. Each scenario describes the technology issue being solved, the environment in which the solution is implemented, and illustrates the implementation using SequeLink middleware.

Scenario 1

Bank International is a large, international banking firm that has grown from multiple acquisitions. It offers many financial services, including general banking, investments, stocks, bonds, and credit cards. Bank International wants to align its business and information technology strategies so that it can better serve current customers as well as new ones. The company also needs to deal with the formidable challenges related to ongoing mergers and acquisitions.

Technology Issues

Bank International wants to move to a network-centric, distributed architecture using Java and CORBA. It will need to integrate legacy applications and migrate quickly to new systems. Distributed transactions are a key element for the future, especially as the company moves toward Internet and business-to-business applications involving multiple data stores.

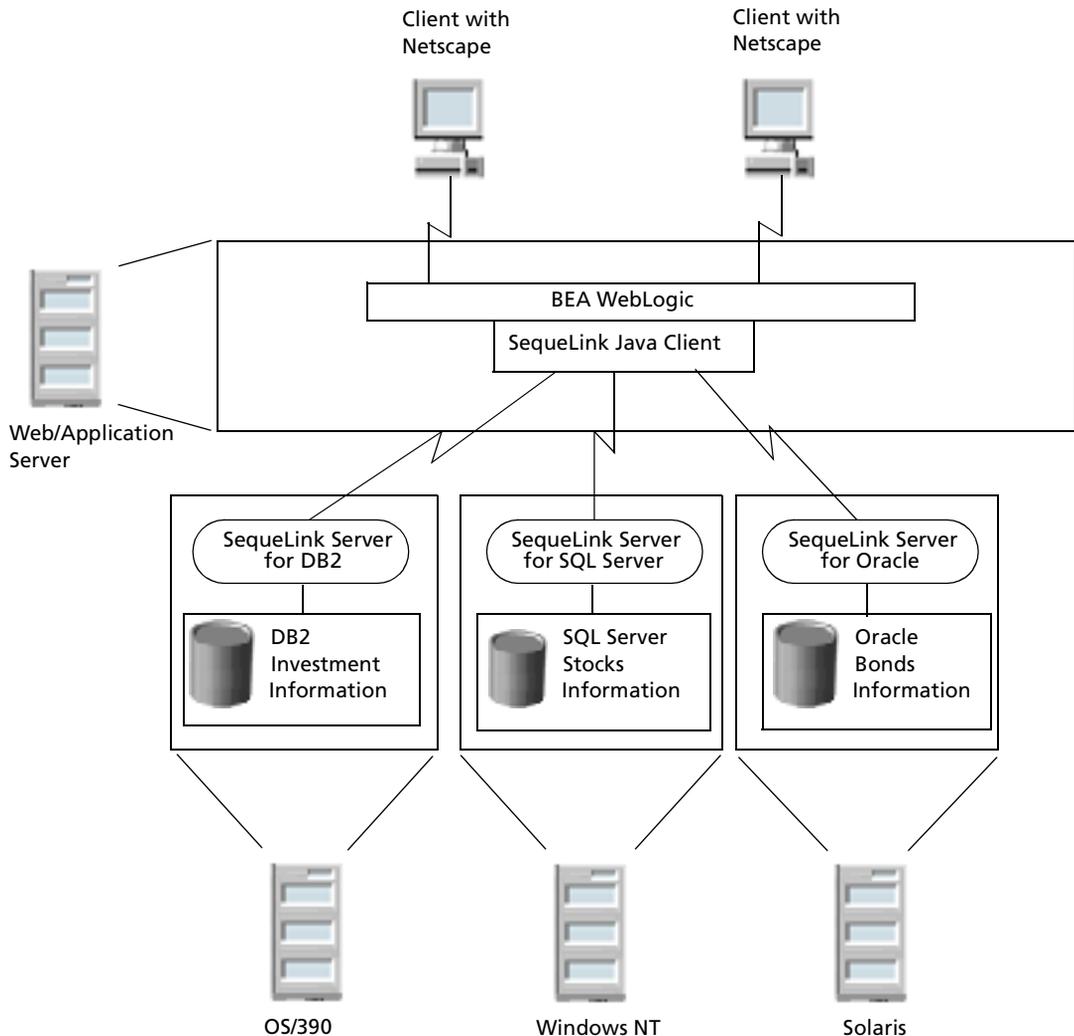
Environment

Because of previous mergers and acquisitions, Bank International has a wide variety of data and systems including DB2 on OS/390, Microsoft SQL Server on Windows NT, and Oracle on Solaris. It uses commercial software, as well as systems developed internally with Inprise jBuilder and some Sun Java development tools. The Web/Application server uses BEA WebLogic.

SequeLink Solution

Because BEA WebLogic is Java-based, it requires the SequeLink Java Client as shown in [Figure 2-1](#).

Figure 2-1. Scenario 1



In this scenario, you install the SequeLink Java Client on the application server, and SequeLink Servers for DB2, Microsoft SQL Server, and Oracle on the database servers. SequeLink allows quick and easy changes to Bank International's environment to accommodate its frequent merger and acquisition activity by allowing administrators to add SequeLink Servers for data access to new data stores and to add different SequeLink Clients for data access applications written in different APIs.

Scenario 2

International Motor Corporation is an automobile manufacturer with several IT implementations. There is no current centralization, but management has decided to control costs through streamlined, thin-client desktop standardization. The company wants to provide better service to remote sites, eventually evolving a system offering a single, company-wide view of all available information.

Technology Issues

Technology issues include the reuse of business logic across applications and the deployment of centralized views of company business systems.

Environment

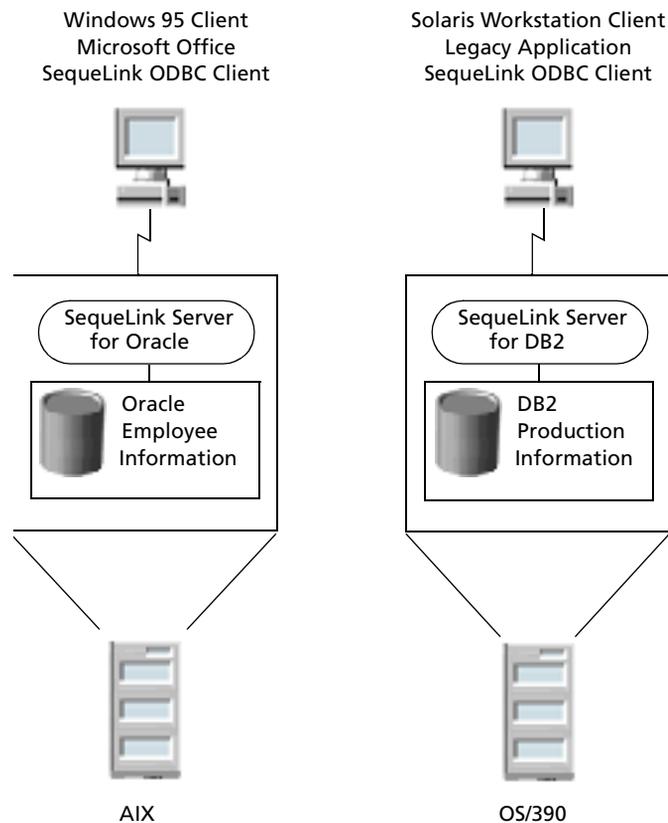
The current environment includes Windows and UNIX clients. The Windows clients are standardized on Microsoft Office applications as well as specific business intelligence and third-party tools. The UNIX clients run legacy UNIX ODBC

applications. Databases are maintained using DB2 on OS/390 and Oracle on AIX.

SequeLink Solution

The client machines use a combination of applications, which run on Windows and UNIX, that require the SequeLink ODBC Client on both platforms as shown in [Figure 2-2](#).

Figure 2-2. Scenario 2



In this example, you install the SequeLink ODBC Client on the client machine and SequeLink Servers for DB2 and Oracle on the database servers.

3 Planning Your SequeLink Configuration

Many of the configuration decisions you need to make will depend on which server and client platforms SequeLink is installed on, which databases you will be accessing, and which SequeLink Client you are using (ODBC, ADO, or Java). See [Chapter 1 “Introduction” on page 17](#) for a description of the SequeLink components and architecture.

This chapter provides information you need to know as you plan your SequeLink configuration, including:

- [“Information You Need Before You Configure” on page 42](#)
- [“Summary of What You Must Configure” on page 42](#)
- [“Planning Your Client Data Sources” on page 43](#)
- [“Planning Your SequeLink ODBC Client and SequeLink ADO Client Configurations” on page 46](#)
- [“Planning SequeLink Data Access Services” on page 47](#)
- [“Planning Your Connection Model” on page 58](#)
- [“Planning Security” on page 59](#)
- [“Planning System Administration” on page 70](#)
- [“Planning Monitoring and Event Tracing” on page 72](#)
- [“Examples: Configuring SequeLink” on page 74](#)

Information You Need Before You Configure

Before you configure your SequeLink environment, you need to know the following information about your SequeLink installation:

- Platforms on which the SequeLink Server software is installed
- Type of SequeLink Servers that are installed—for example, SequeLink Server for Oracle
- Platforms on which the SequeLink Client software is installed
- Types of SequeLink Clients that are installed—ODBC, ADO, or Java

Summary of What You Must Configure

This section describes the items that you are required to configure for SequeLink.

After you have completed the installation of the SequeLink Client and SequeLink Server as documented in the *SequeLink Installation Guide*, you must configure a client component. The client component you configure depends on your SequeLink Client:

- For a SequeLink ODBC Client or SequeLink ADO Client, you must configure a client data source.
- For a SequeLink Java Client, you must configure a client data source or a connection URL.

A client data source is a minimal data source that is stored on the client. It provides connection information to a specific server. By default, a client data source uses the default server data source

for the characteristics of the connection between the SequeLink Client and the database.

This default server data source, named *Default*, is automatically created and configured when you install the SequeLink Server. A server data source defines the characteristics of the connection between the SequeLink Client and the database. You can modify the definition of the default server data source if needed. See [“Planning SequeLink Data Access Services” on page 47](#) for information about the default behavior of the default server data source.

NOTE: A server data source, although it resides on the server, also contains settings that are relevant to the SequeLink Client. Centralizing data source information on the server allows SequeLink to provide "near-zero" client administration.

Planning Your Client Data Sources

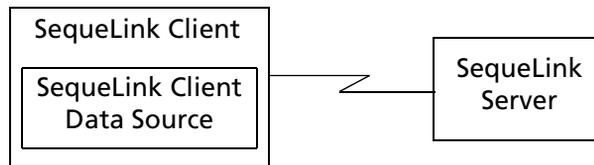
Configuring SequeLink Client data sources is a simple task, but, before you configure the data sources, you must make the following decisions:

- You must decide which server connection method to use—direct connection or retrieve the connection parameters from a Lightweight Directory Access Protocol (LDAP) directory.
- You must decide whether to reference the default server data source or another server data source.

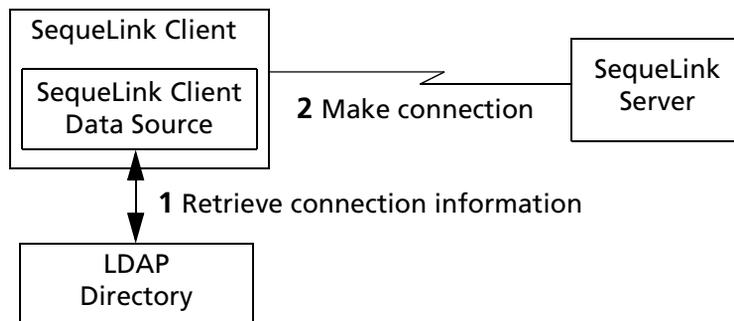
Configuring a Connection to a SequeLink Server

SequeLink allows you to configure a connection from the SequeLink Client to a SequeLink Server in either of the following ways:

- Specify the connection information directly in the SequeLink Client data source. This information includes the TCP/IP address (or host name) and port of the SequeLink Server.



- Specify a Distinguished Name (DN) identifying a specific entry in an LDAP directory. The LDAP directory entry contains the information needed to connect to the SequeLink Server. Using this method, the SequeLink Client retrieves the connection information from the LDAP directory.



The advantage of using LDAP is that you can centrally store connection information, which provides the flexibility to make environment changes. For example, if you move the database and SequeLink Server to a different server, you do not have to reconfigure your user applications or multiple SequeLink Client data sources that access the SequeLink Server; the connection

information is specified in an LDAP directory, not in the SequeLink Client data source. Therefore, you need only make a single change in the LDAP directory entries.

For SequeLink Java Clients, the JDBC client data source stores connection instructions in a JNDI infrastructure, which can support LDAP.

Referencing a Server Data Source

If you do not specify a server data source when configuring a client data source, the client data source uses the default server data source. The behavior of the default server data source is documented in [“General Data Source Default Behavior” on page 50](#). If the default server data source definition does not meet your needs, you can modify the default server data source, or you can create a new server data source and then specify the new server data source when configuring your client data source.

Planning Your SequeLink ODBC Client and SequeLink ADO Client Configurations



On Windows platforms, ensuring that multiple SequeLink ODBC or ADO Clients have the same configuration is an easy task when you create a predefined client installation image called a *Quick Install image*. After you define a Quick Install image, users can install this image on their client machines.

Configuring Quick Install Images

A Quick Install image can be stored on a file server and used for all SequeLink ODBC or ADO Client installations within a workgroup. Installing a Quick Install image requires minimal user interaction and ensures that every client for which you are responsible has the same configuration. You can also define multiple Quick Install images, which allows you to customize each image for different workgroups within your organization.

When configuring a Quick Install image, you define the following information:

- Whether the installation is a workstation or network installation
- Location of the installation directory
- Which SequeLink components to install
- Which mode the installation will run in (interactive or batch)
- Whether client data sources are installed (see the following section [“Including Client Data Source Configurations in Quick Install Images”](#))

Refer to the *SequeLink Installation Guide* for instructions on how to define a Quick Install image.

Including Client Data Source Configurations in Quick Install Images

As part of a Quick Install image, you can specify client data sources that are to be installed on the client machines, which ensures that each client has the same client data source configuration. See [“Planning Your Client Data Sources” on page 43](#). The client data sources that are part of a Quick Install image must first be exported to a data source file using the SequeLink Data Source SyncTool. The SequeLink Data Source SyncTool allows you to create data source files and export data source definitions to data source files. See the *SequeLink Administrator’s Guide* for information about how to use the SyncTool.

To maintain client data source configurations, you can distribute a data source file and have users export the data source definitions from the data source file to their client machines. This ensures that the same client data source configuration is installed on all client machines.

Planning SequeLink Data Access Services

A SequeLink data access service, which resides on the server, allows a SequeLink Client to connect to a database and is required for your SequeLink configuration. On Windows NT, Windows 2000, and UNIX, a SequeLink data access service is automatically created and configured when you install SequeLink Server, so you do not have to define one. On OS/390, defining a SequeLink data access service is part of the installation process as described in the *SequeLink Installation Guide*.

Part of the definition of each data access service is a server data source. A data access service has one default server data source

defined. You can define additional server data sources for a data access service, but a SequeLink Client data source can reference only one server data source. Therefore, if you create a second server data source for a data access service and you want the SequeLink Client data source to reference the second server data source, you must create a new client data source and configure it to reference the new server data source.

This section provides the definition of the default (or newly created) data access service on Windows NT, Windows 2000, UNIX, and OS/390 so that you can decide whether the default definition fits your needs. If it does not, you can change the values of the service's attributes. The attributes can be modified using one of the SequeLink Manager tools. For instructions on using the SequeLink Manager, refer to the *SequeLink Administrator's Guide*.

This section also provides examples of why you would change the default definition and why you may want to have multiple server data sources defined for a data access service.



Default Behavior on Windows NT, Windows 2000, and UNIX

The following sections define the default behavior of the default SequeLink data access service and provide the name of the attribute that dictates the behavior in parentheses. The attributes that begin with "DataSource" (for example, DataSourceReadOnly) are the attributes that define a server data source. For complete information about service attributes, see the *SequeLink Administrator's Guide*.

Connection Model Default Behavior

- The connection model used is ThreadPool. See [“Planning Your Connection Model” on page 58](#). (ServiceConnectionModel)
- The number of prestarted threads in the thread pool is 8. (ServiceMinThreads)
- The maximum number of threads to which the thread pool can increase to accommodate peak activity is 64. (ServiceMaxThreads)
- The time allowed for thread-pool synchronization actions to take place before an internal error is generated is 60000 milliseconds. (ServiceInternalTimeout)
- The idle time after which a thread allocated to a specific connection is released to the thread pool is 2000 milliseconds. (DataSourceThreadRpcTimeOut)
- The maximum number of requests after which a thread allocated to a specific connection is released to the thread pool is 10. (DataSourceThreadMaxRpc)

Event Handling Default Behavior

- The size of the event trace file is 1000000 bytes. (ServiceEventTraceSize)
- The default location of the event trace file is the tracing subdirectory of the SequeLink Server installation directory. The default location is
 \Program Files\Merant\SLServer51\tracing on Windows NT and Windows 2000 and /usr/slserver51/tracing on UNIX. (ServiceEventTraceLocation)
- The time to wait for a data access service response before timing out is 1000 milliseconds. (ServiceEvQPingTimeout)

See [“Planning Monitoring and Event Tracing” on page 72](#).

Security Default Behavior

See [“Planning Security on Windows and UNIX”](#) on page 67 for details.

Debug Default Behavior

Debugging is disabled. (ServiceDebugLogLevel and ServiceDebugLogPath)

Data Privacy Default Behavior

Cleartext messages are transmitted between the client and server over the network; no data scrambling is used. (ServiceEncryptionAlgorithm)

NOTE: User IDs and passwords are *never* sent as cleartext.

General Service Default Behavior

- The service tries to recover from unexpected exceptions, rather than passing the exception to the operating system, which would result in a core dump. (ServiceCatchExceptions)
- The SequeLink service supports only standard ASCII/EBCDIC transliteration. (ServiceCodePage)

General Data Source Default Behavior

- Server data sources have no name or description specified. (DataSourceName and DataSourceDescription)
- No schema names are specified to filter the result set returned by SQLTables and SQLProcedures (ODBC), getTables and getProcedures (JDBC), and TABLES and PROCEDURES (OLE DB/ADO). (DataSourceSchemaFilterList)

- No list of tables types is specified to filter the result set returned for SQLTables (ODBC), getTables (JDBC), and TABLES (OLE DB/ADO). (DataSourceTableTypeFilterList)
- The size of the buffer to use for array fetch is 65536 bytes. (DataSourceArrayFetchMaxBytes)
- The SequeLink Client can perform all supported SQL statements; the database connection is **not** read-only. (DataSourceReadOnly)
- No default database catalog is used when connected to a SequeLink data access service. (DataSourceCurrentCatalog)

DB2 Service Default Behavior

- The default database catalog is an empty string, which causes the SequeLink ODBC Client and SequeLink ADO Client to prompt for a valid DB2 database alias when connecting to DB2 V6 or DB2 V7. (DataSourceCurrentCatalog)
- No DB2 connection options are set for the service. (DataSourceDB2ConnectOptions)

Informix Service Default Behavior

- The default database catalog is an empty string, which causes the SequeLink ODBC Client and SequeLink ADO Client to prompt for a valid Informix database name. (DataSourceCurrentCatalog)
- A value for the Informix environment variable INFORMIXSERVER is set for the service. (ServiceEnvironmentVariable on UNIX, DataSourceINFormixServer on Windows NT/2000)
- A value for the Informix environment variable INFORMIXDIR is set for the service. (ServiceEnvironmentVariable on UNIX, DataSourceINFormixDir on Windows NT/2000)

- Informix delimited identifiers are enabled for the service. (ServiceEnvironmentVariable on UNIX, DataSourceINFIformixDelimident on Windows NT/2000)

Microsoft SQL Server Service Default Behavior

- No default catalog is specified for use with Microsoft SQL Server when connecting to the database. (DataSourceCurrentCatalog)
- No warning is passed to the client application if Microsoft SQL Server encounters a problem opening server-side cursors. (DataSourceMSSCursorWarnings). The type of cursor used is specified using the DataSourceMSSCursorType attribute.
- SequeLink Server uses a server-side cursor, which allows multiple concurrent statements (and cursors) to be active at the same time for each connection. (DataSourceMSCursorType)

Oracle Service Default Behavior

- Synonyms of remote Oracle objects are **not** supported in catalog statements. (DataSourceORASynDBLinkObjSupp)
- All Oracle API calls that must be executed to establish a connection with the Oracle database will be serialized. (ServiceORASerializeLogon)

Sybase Service Default Behavior

A value for the Sybase network address (*hostname, port*) is set for the service. (DataSourceSYBNetworkAddress)

Default Behavior on OS/390

The following sections define the default behavior of a SequeLink data access service and provides the name of the attribute that dictates the behavior in parentheses. The attributes that begin with DataSource (for example, DataSourceReadOnly) are the attributes that define a server data source. For complete information about service attributes, see the *SequeLink Administrator's Guide*.

Connection Model Default Behavior

The connection model used is Process/Connection. See [“Planning Your Connection Model” on page 58](#). (ServiceConnectionModel)

Event Handling Default Behavior

The size of the event trace file is 1000000 bytes. (ServiceEventTraceSize)

See [“Planning Monitoring and Event Tracing” on page 72](#).

Security Default Behavior

See [“Planning Security on OS/390” on page 68](#) for details.

Debug Default Behavior

Debugging is disabled. (ServiceDebugLogLevel and ServiceDebugLogPath)

Data Privacy Default Behavior

Cleartext messages are transmitted between the client and server over the network; no data scrambling is used.

(ServiceEncryptionAlgorithm)

NOTE: User IDs and passwords are *never* sent as cleartext.

General Service Default Behavior

- The name of the service program is VAICDB2S and cannot be modified, except on request of MERANT technical support. (MVSServiceLoadModule)
- The SequeLink service supports only standard ASCII/EBCDIC transliteration. (ServiceCodePage)
- When the data store does not run in a specific code page, the code page is read from the setting of the MVSServiceCodePageNr attribute. The default is 37. (MVSServiceCodePageNr)
- The maximum number of concurrent sessions for a service is 200. When this threshold is reached, the server refuses subsequent connection requests for the service. (MVSServiceMaximumSession)

General Data Source Default Behavior

- Server data sources have no name or description specified. (DataSourceName and DataSourceDescription)
- Transaction isolation levels are supported on connections. The default is Read Committed (cursor stability). (DataSourceTransactionIsolation)
- Cursors are not destroyed when the transaction is committed. (DataSourceCursorHold)

- No schema names are specified to filter the result set returned by SQLTables and SQLProcedures (ODBC), getTables and getProcedures (JDBC), and TABLES and PROCEDURES (OLE DB/ADO). (DataSourceSchemaFilterList)
- No list of table types is specified to filter the result set returned for SQLTables (ODBC), getTables (JDBC), and TABLES (OLE DB/ADO). (DataSourceTableTypeFilterList)
- The SequeLink Client can perform all supported SQL statements; the database connection is **not** read-only. (DataSourceReadOnly)
- No database catalog is used when connected to a SequeLink data access service (DataSourceCurrentCatalog)

DB2 Service Default Behavior

- The default DB2 plan name for the server data source is SWDB2PL. This plan contains packages bound with different transaction isolation levels. The name of these packages consists of a collection name and a suffix that indicates the transaction level of the package. (MVSDDataSourceDB2Plan and DataSourceDB2CollectionPrefix)
- No default catalog is specified for use with DB2 when connecting to the database. (DataSourceCurrentCatalog)
- Database cursors are opened without the HOLD attribute for DB2. (DataSourceCursorHold)

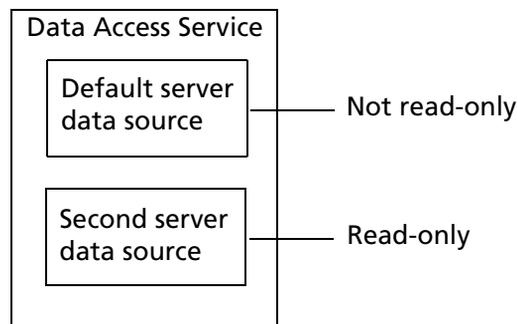
Examples of Changing Default Behavior

The examples in this section demonstrate possible scenarios and reasons that may have to modify the default behavior of a data access service. In addition, it describes reasons you may need more than one server data source for a data access service.

Example A:

You may want most of your client applications that use a specific data access service (such as SequeLink Server for Oracle8) to connect to the database with read-only functionality. Therefore, you would change the value of the `DataSourceReadOnly` service attribute to read-only.

The `DataSourceReadOnly` service attribute is an attribute of a server data source. Because changing the value of the default server data source would cause all connections to the database to be read-only, you would choose to create a second server data source for the data access service and change the value of the `DataSourceReadOnly` service attribute of the new server data source as shown:



Example B:

You may want to configure your data access service to use data scrambling. Remember that the default is that cleartext messages are transmitted between the client and server over the network.

In this case, you would change the value of the `ServiceEncryptionAlgorithm` service attribute to one of the three available data scrambling options: byte swapping, DES, or 3DES. The `ServiceEncryptionAlgorithm` attribute is set at the data access service level and is not part of the definition of a server data source. Therefore, any client application that uses the data access service will have data scrambling set.

Example C:

SequeLink ODBC Clients support multiple workarounds to circumvent limitations in some applications (such as Microsoft Access). Using SequeLink, these workarounds can be configured at the server for a specific server data source. This example shows why you may want to configure multiple server data sources for a data access service.

The service attribute `DataSourceFetchTimeStampAsString` specifies whether a workaround for a Microsoft Access problem with timestamps is turned on. The workaround is turned off by default (`FALSE`). You can create a new server data source and set the `DataSourceFetchTimeStampAsString` attribute to `TRUE` to turn on this workaround. With the workaround turned on, the Microsoft Access application can use the client data source that references this server data source.

Planning Your Connection Model

The type of connection model you choose for your SequeLink configuration partly depends on your SequeLink Server platform and whether you are using the distributed transaction functionality of a DBMS. SequeLink provides the following types of connection models:

- The **ThreadPool** connection model starts SequeLink with a preallocated minimum number and maximum number of threads that can be shared by multiple SequeLink Clients connected to the SequeLink Server. It provides optimum scalability—many client connections can be serviced with the same system resources on the server. It is the default connection model for all platforms except OS/390; the ThreadPool connection model cannot be used on OS/390.



NOTE: If you are using distributed transactions with DB2 Universal Database on Windows or UNIX platforms, do not use the ThreadPool connection model; use the Process/Connection model.

- The **Process/Connection** connection model creates a separate operating system task for each SequeLink Client connection request. This is the only connection model supported on OS/390.



NOTE: If you are using distributed transactions with DB2 Universal Database on Windows or UNIX platforms, use the Process/Connection model.

- The **Thread/Connection** connection model provides a dedicated thread for each SequeLink Client connection to a SequeLink Server. Use the Thread/Connection connection model for client applications that are database-intensive, such as bulk load or bulk transfer applications.

Planning Security

This section first discusses the security mechanisms supported by SequeLink and identifies the service attributes that must be set to configure each security mechanism. Next, two planning sections are provided—one for Windows and UNIX, and another for OS/390—that discuss the default behavior of security on each platform.

About SequeLink Security

SequeLink security supports security mechanisms for the following purposes:

- Verification of a user by the SequeLink Server. The **Authentication** security mechanism allows the SequeLink Server to verify the identity of the user.
- Types of requests that are accepted by the server. The **Authorization** security mechanism controls whether the user can send data access requests and administrative (SequeLink Manager) requests, and whether the requests can be accepted by the server.
- Connection to a data store using the following security mechanisms:
 - **Data Store Logon** controls whether a user who is connected to the SequeLink Server can connect to the data store.
 - **Application IDs** control whether a client application can connect to the data store. This mechanism adds a layer of security on top of Data Store Logon.
- Types of SQL statements accepted by the data store. The **ReadOnly** security mechanism controls whether the data store connection is read-only.

- The privacy of the data being transmitted. The **data privacy** security mechanism ensures that data transmitted between the client and server is kept private using data scrambling methods and encryption through Secure Socket Layer (SSL) (SSL is supported for Java environments only).

Authentication

Authentication allows the SequeLink Server to verify the identity of the SequeLink Client when the client connects to the SequeLink Server. If authentication fails, the SequeLink Client will disconnect from the server. Authentication must be set separately for users (people who send data access requests) and administrators (people who send SequeLink Manager requests).

Even though a user may be able to connect to the SequeLink Server, this does not mean that the user automatically has access to the database that the SequeLink Server services. Access to the database is controlled by Data Store Logon (see [“Data Store Logon” on page 64](#)) and Application IDs (see [“Application IDs” on page 64](#)).

Depending on the combination of client and server platforms involved in the connection, SequeLink supports the following authentication mechanisms:

- **Anonymous.** The SequeLink Server accepts connections from any SequeLink Client without verifying the client’s identity. This is the default on Windows and UNIX for users (ServiceAuthMethods=anonymous).
- **Integrated NT.** This option is supported for connections between SequeLink ODBC Clients and SequeLink ADO Clients on Windows and SequeLink Server for Windows NT servers only. The SequeLink Server verifies the identity of the SequeLink Client using the client’s Windows network logon credentials instead of a Windows user and password.



- **Operating system user ID and password.** The SequeLink Server verifies the identity of the SequeLink Client using a user ID and password that must be valid for the platform on which the SequeLink Server is running. If verified, the server accepts the user ID as the identity of the client and permits the connection. This is the default on Windows and UNIX for administrators (ServiceAdminAuthMethods=OSlogon(UID,PWD)).

Authorization

After the SequeLink Server has authenticated the client, SequeLink verifies that the client is authorized to perform data access activities or SequeLink Manager activities. SequeLink supports authorization for data access requests and for SequeLink Manager requests. You configure the authorization for the two types of requests separately. Authorization options depend on your SequeLink Server platform.

Authorization for Windows and UNIX

You configure the authorization for data access requests and for SequeLink Manager requests separately:

- To configure authorization for data access, set the ServiceUser attribute. If you want to configure authorization for user groups defined on Windows NT or UNIX, you can set the ServiceGroupUser attribute. These attributes should be added to data access services only.
- To configure authorization for SequeLink Manager requests, set the ServiceAdministrator attribute. If you want to configure authorization for user groups defined on Windows NT or UNIX, you can use the ServiceGroupAdministrator attribute. These attributes should be added to SequeLink Agent services only.

The `ServiceUser` attribute and `ServiceAdministrator` attributes can have the following values:

- **Everyone.** The SequeLink Server will process all requests sent by the user, regardless of how the user is authenticated. For example:

```
ServiceUser=everyone
```

If you set authentication to anonymous, you must set authorization to everyone (`ServiceUser=everyone` or `ServiceAdministrator=everyone`).

This is the default for `ServiceUser`.

- **Authenticated.** The SequeLink Server will process all requests sent by the user if the user can be authenticated (authentication is set by the `ServiceAuthMethods` and `ServiceAdminAuthMethods` attributes). For example:

```
ServiceAdministrator=authenticated
```

- **User_id.** The SequeLink Server will process all requests sent by the user if the user ID has been specified as authorized. For example, to configure permission for marym to send data access requests, you would set the following attribute for the data access service:

```
ServiceUser=marym
```

And, to configure permission for Marym to send SequeLink Manager requests, you would set the following attribute for the SequeLink Agent service:

```
ServiceAdministrator=marym
```

This is the default for `ServiceAdministrator`. You specify a user ID as the default administrator ID during the installation of the SequeLink Server.

NOTES:

- Alternatively, you can set the `ServiceUserGroup` and `ServiceAdministratorGroup` attributes to configure authorization for groups of users defined on Windows NT or UNIX.
- On Windows NT, users who are allowed to manage SequeLink services using the SequeLink Manager must have NT administrator rights.

Authorization for OS/390

On OS/390, you can configure authentication with or without additional authorization for both SequeLink data access services and SequeLink Agent services. The attributes that control this functionality are `MVSServiceSecurity` for SequeLink data access services and `MVSServiceAdminSecurity` for SequeLink Agent services.

To configure authorization on OS/390, you must specify a security class and a security resource in this security class by setting the following attributes:

- `MVSServiceSecurityResource` and `MVSServiceSecurityClass` for SequeLink data access services.
- `MVSServiceAdminSecurityResource` and `MVSServiceAdminSecurityClass` for SequeLink Agent services.

Data Store Logon

Once a connection is established, authentication is complete, and the type of requests accepted by the server has been established, a connection from the SequeLink Server to the database can be established by using either of the following methods:

- Specifying data store logon information (a valid DBMS user ID and password). This is the default for Windows and UNIX (`DataSourceLogonMethod=DBMSLogon(UID, PWD)`).
- Allowing the database to inherit the logon user ID that was established during the authentication process. This method **must** be used for OS/390, but it also can be used for Windows and UNIX (`DataSourceLogonMethod=OSIntegrated`).

Application IDs

Application IDs are alphanumeric strings passed by a SequeLink Client that identify the client application to a SequeLink service that has been configured to accept connections only from specific application IDs.

Application IDs add another layer of security for the connection to the data store beyond that provided by the Data Store Logon security mechanism. Data Store Logon allows all users of client applications to access the data store if the users meet the qualifications set by Data Store Logon. Using application IDs, you can restrict connections to the data store to only those client applications that identify themselves to the SequeLink Server through an application ID.

ReadOnly

SequeLink allows you to configure the types of SQL statements the data store connection will accept:

- Select statements only, which makes the connection read-only
- Select statements and Stored Procedures
- All SQL statements

The service attribute that controls this functionality is `DataSourceReadOnly`.

Data Privacy

SequeLink provides data scrambling (for all supported environments) and real encryption through the use of SSL (SequeLink Java environments only) to ensure the privacy of data.

Data Scrambling

Data scrambling ensures that no cleartext messages are transmitted between the client and server over the network. SequeLink provides the following implementations of data scrambling:

- **Fixed-key DES** operates using a 56-bit key.
- **Fixed-key 3DES** operates using a 168-bit key.
- **Byte swapping** means that bytes of data are randomly swapped to scramble data. Different encoded mappings are used for different sessions.

Data scrambling does not offer the same level of security as SSL and its use may degrade performance. Data scrambling is not enabled by default.

NOTE: Even if you choose not to use a data scrambling method, user IDs and passwords are *never* sent as cleartext.

To configure SequeLink to use DES, 3DES, or byteswap, you must set the `ServiceEncryptionAlgorithm` service attribute, for example, `ServiceEncryptionAlgorithm=DES`. The default is none, which means cleartext messages are transmitted between the client and server over the network.

SSL (Java Environments Only)

Secure Socket Layer (SSL) encryption provides data encryption, server authentication, and message integrity for TCP/IP connections using the following methods:

- **Asymmetric cryptographic algorithms** protect the exchange of symmetric encryption keys. SequeLink supports the following asymmetric cryptographic algorithm classes:
 - **Anonymous.** The exchange of the symmetric key for the data transfer is protected by an asymmetric key agreement protocol, but the client does not verify the identity of the server. The anonymous mechanism provides protection against passive eavesdropping on communication lines, preventing someone who is monitoring network traffic from deciphering the exchanged data. It does not provide protection from “man-in-the-middle” security infiltrations, in which intruders position themselves between the client and the server, pretending to the client to be the server and vice versa, and allowing the intruders to intercept, inspect, and possibly alter the data exchanged between the client and the server.
 - **Server authentication.** When communication begins, the server identifies itself to the client, using a digital certificate as proof of identity. The client verifies this certificate to ensure that the server is really the party with whom it wants to communicate.

- **Symmetric cryptographic algorithms** encrypt and decrypt the actual data.
- **Message digest algorithms** ensure message integrity.

The combination of all these algorithms is called a *cipher suite*. The SSL standard defines the cipher suites that can be specified. The actual availability of a cipher suite is determined by the underlying implementation. For more information about configuring SSL and specifying cipher suites, refer to the *SequeLink Administrator's Guide*.

NOTE: SequeLink supports Netscape's Transport Layer Security (TLS) 1.0 through its SSL implementation.

Planning Security on Windows and UNIX

This section lists the default security behavior for users (people who will be sending data access requests to the server) and administrators (people who will be sending SequeLink Manager requests to the server).

Default Security Behavior for Users

The SequeLink Server accepts connections from all users, but only the users who can provide a valid DBMS user ID and password are allowed to access the database. The database connection accepts all types of SQL statements. Once connected to the database, the database security system guarantees that the user can only perform actions that are allowed by the database administrator. Messages (except for user IDs and passwords) sent between the client and the server are sent as cleartext. The service attributes are set as follows:

```
ServiceAuthMethods=anonymous  
ServiceUser=everyone
```

```
DataSourceLogonMethod=DBMSLogon (UID, PWD)
DataSourceReadOnly=No
ServiceEncryptionAlgorithm=none
```

Default Security Behavior for Administrators

Only the person who logs on using the administrator ID entered when the SequeLink Server software was installed is allowed to manage the SequeLink environment. The SequeLink Server Setup prompts for a user ID for the SequeLink administrator when you install the SequeLink Server.

```
ServiceAdminAuthMethods=OSlogon (UID, PWD)
ServiceAdministrator=User_ID
```

NOTE: On Windows NT, the SequeLink administrator must have NT administrator rights.

Refer to the *SequeLink Administrator's Guide* for complete information about configuring SequeLink security.

Planning Security on OS/390

This section lists the default security behavior for users (people who will be sending data access requests to the server) and administrators (people who will be sending SequeLink Manager requests to the server).

Default Behavior for Users

The SequeLink Server accepts connections from all users, but only users who can provide a valid operating system user ID and password are allowed to access the database. The database connection accepts all types of SQL statements. Once connected to the database, the database security system guarantees that the user can only perform actions that are allowed by the database administrator. Messages (except for user IDs and

passwords) sent between the client and the server are sent as cleartext. The service attributes are set as follows:

```
MVSServiceSecurity=SAFBASIC  
ServiceAuthMethods=OSLogon (UID, PWD, NWPD)  
DataSourceLogonMethod=OSIntegrated  
DataSourceReadOnly=No  
ServiceEncryptionAlgorithm=none
```

Default Behavior for Administrators

The SequeLink Server accepts connections from all users, but only users who can provide a valid operating system user ID and password are allowed to access the SequeLink Agent service, which means they can manage the SequeLink environment. The service attributes are set as follows:

```
MVSServiceAdminSecurity=SAFBASIC  
ServiceAdminAuthMethods=OSLogon (UID, PWD)
```

Refer to the *SequeLink Administrator's Guide* for complete information about configuring SequeLink security.

Planning System Administration

SequeLink provides the following options for configuring and managing your SequeLink environment:

- **Local system administration** allows you to configure and manage your SequeLink environment using the SequeLink Manager installed locally on a SequeLink Server.
- **Remote system administration** allows you to configure and manage your SequeLink environment using the SequeLink Manager installed on the desktop of a networked client.

NOTE: Only SequeLink 5.1 services can be configured, managed, or monitored with the SequeLink Manager 5.1.

Local System Administration

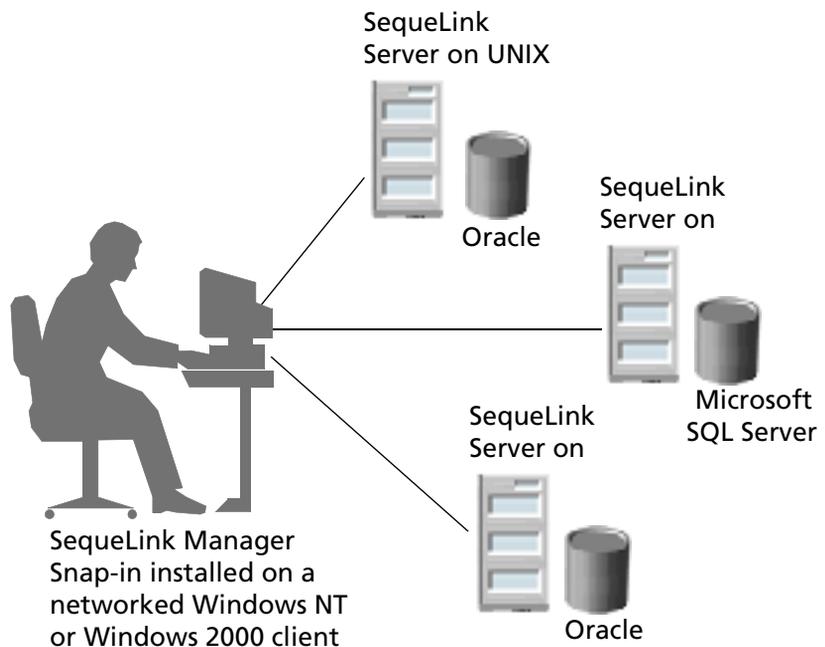
You can use the SequeLink Manager locally from the SequeLink Server to configure and manage your SequeLink environment; however, the SequeLink Manager tool you can use locally depends on the platform of the SequeLink Server. See [“Understanding the SequeLink Manager Tools” on page 25](#) for a description of the SequeLink Manager tools and information about the platforms on which they can be installed.

Although local system administration (administration from the server) works for all SequeLink Server platforms, it is not always the most convenient way to handle administration of your SequeLink environment. It may be, however, the best solution for OS/390 because remote administration for OS/390 is limited. To configure and manage SequeLink services on OS/390, or create OS/390 core entities such as DB2 interfaces, you must use the SequeLink Manager Tool for OS/390 locally.

Remote System Administration

Remote system administration allows you to configure and manage your data access environment from the convenience of your desktop regardless of your SequeLink Server platform. For example, suppose you are responsible for administering an environment with distributed data access involving a variety of data stores across your enterprise, such as Oracle on UNIX and Windows, and Microsoft SQL Server on Windows as shown in [Figure 3-1](#). You can install the SequeLink Manager Snap-in on a Windows NT or Windows 2000 networked client and perform administration tasks, such as configuring SequeLink service settings, from the convenience of your desktop.

Figure 3-1. Remote System Administration for Data Access Environments



If you prefer to use a command-line tool rather than a GUI tool, you can use the SequeLink Manager Command-Line Tool to remotely administer your SequeLink environment.

NOTE: The availability of remote system administration depends on your SequeLink license.

Planning Monitoring and Event Tracing

SequeLink *monitoring* allows the SequeLink administrator to see what is currently happening in the SequeLink environment. SequeLink *event tracing* allows the SequeLink administrator to store information about events that occur in an event trace file, allowing them, for example, to check on events that happened overnight.

To enable monitoring and event tracing for SequeLink services, you must have a monitoring profile and an event trace profile configured for each service you want to monitor. On Windows NT and UNIX, both a monitoring and an event trace profile are enabled when you install SequeLink Server.

Monitoring

SequeLink provides the following levels of monitoring for both SequeLink Agent and SequeLink data access services, listed here from highest-level to lowest-level:

- **Service monitoring** monitors these activities by service:
 - Statistics of received packets and sent packets
 - Sessions started and statements opened
 - Active statements and sessions
 - Fetched rows and affected rows
 - Transactions

- **Session monitoring** monitors these activities by session within a service:
 - Statistics of received packets and sent packets
 - Statements opened and active statements
 - Fetched rows and affected rows
 - Transactions
 - Information about each session, such as start time, client information (network address, data source used by the client, and type of client), native database session identification, and database user
- **Statement monitoring** monitors these activities by statement within a session:
 - Fetched rows and affected rows
 - SQL statements issued

Event Tracing

Event tracing provides a method for the SequeLink administrator to store persisted information about an event. Events are generated when a client application accesses data and when certain server activities take place (such as when a service starts and an error occurs). SequeLink can trace the following type of events:

- | | |
|-----------|---------------|
| ■ Service | ■ Statement |
| ■ Session | ■ Transaction |
| ■ Network | ■ Others |
| ■ Error | |

Examples: Configuring SequeLink

In this example, the SequeLink ODBC Client has been installed on Windows NT workstations and the SequeLink Server for Oracle has been installed on a Windows NT Server. Most client applications will connect to the Oracle8 database with read-only functionality, but one of the client applications must have update capability.

Four user applications use the SequeLink Server for Oracle to connect to an Oracle8 database. One application is an order-entry application that needs the ability to update the database. The other applications are reporting applications that must **not** have the ability to update; therefore, they must connect with read-only functionality to the database.

We have decided to place an added layer of security on the order-entry application because it allows updates. The security mechanism that will be used is application IDs (see [“Planning Security” on page 59](#)). By using the application ID security mechanism, we can ensure that if one of the reporting applications is modified to use the same client data source as the order-entry application, the reporting application will be unable to update the database because it is not configured with an application ID.

To configure SequeLink for use with the report applications, we will modify the default server data source to make it read-only by setting the `DataSourceReadOnly` attribute to `Select`, which means only SQL `Select` statements are allowed. Also, we will create a client data source that uses the default data source. We will name the client data source `OR8ReadOnly`. The reporting applications will connect to the Oracle8 database using the `OR8ReadOnly` client data source.

Also, we will configure SequeLink for use with the order-entry application. To make sure that the application is certified to the server, we will assign it an application ID, and set the

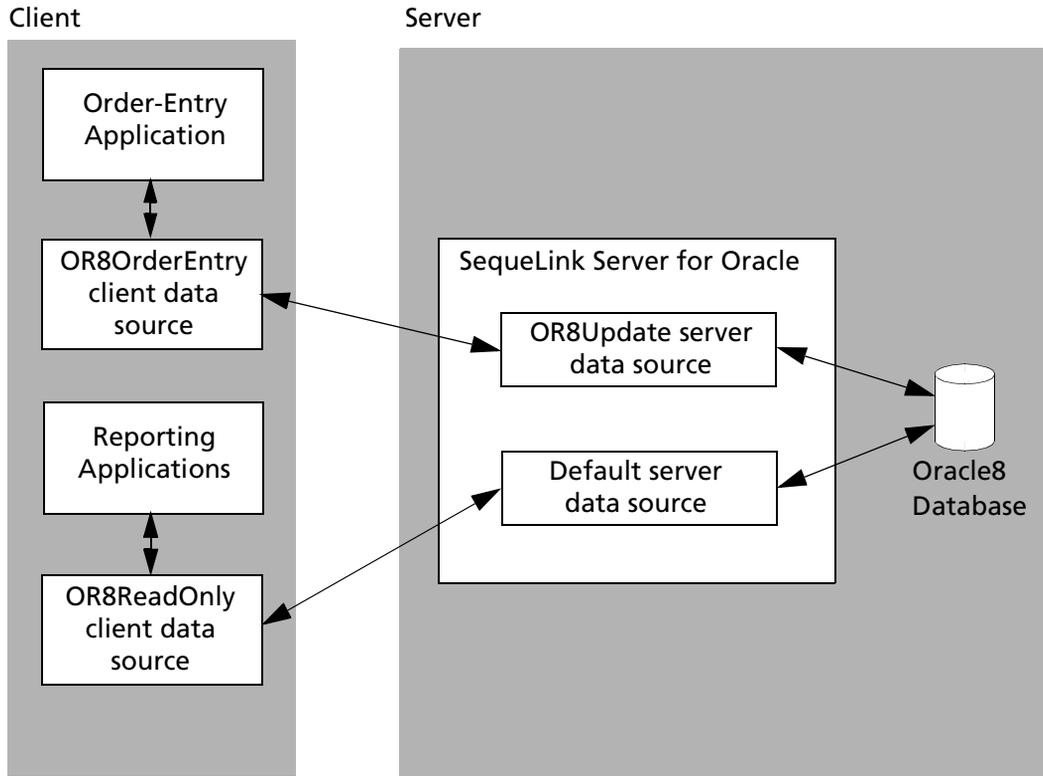
DataSourceApplID service attribute to the value of the order-entry application's ID. We must create and configure a server data source that sets the DataSourceApplID attribute; we will name the server data source Or8Update. In addition, we must create and configure a client data source that references the Or8Update server data source. We will name the client data source Or8OrderEntry.

The following summary lists the components that must be configured:

- A new server data source named OR8Update must be created. This server data source uses application IDs to certify the order-entry application to the server.
- The default server data source must be modified to make it read-only.
- An ODBC client data source named OR8OrderEntry must be created. This client data source will reference the newly created server data source, OR8Update. The order-entry application must connect to this ODBC client data source. For information about how to connect ODBC applications to SequeLink Client data sources, refer to the *SequeLink Administrator's Guide*.
- An ODBC client data source named OR8ReadOnly must be created. This client data source will reference the default server data source. The reporting applications must connect to this ODBC client data source.

Figure 3-2 shows how all of the components fit together.

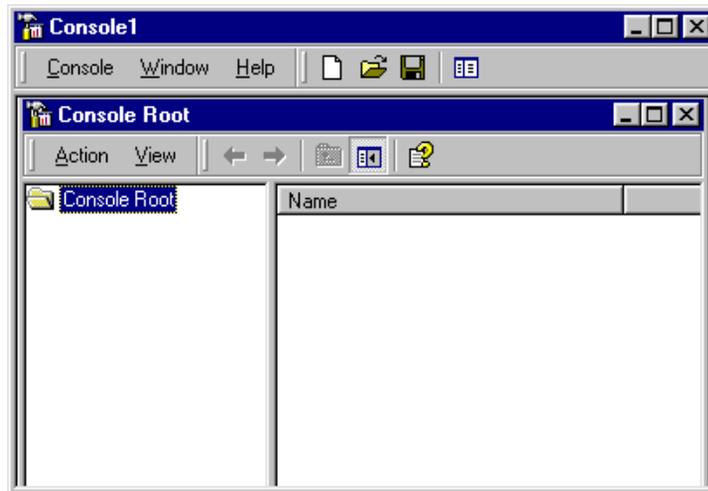
Figure 3-2. Example Configuration



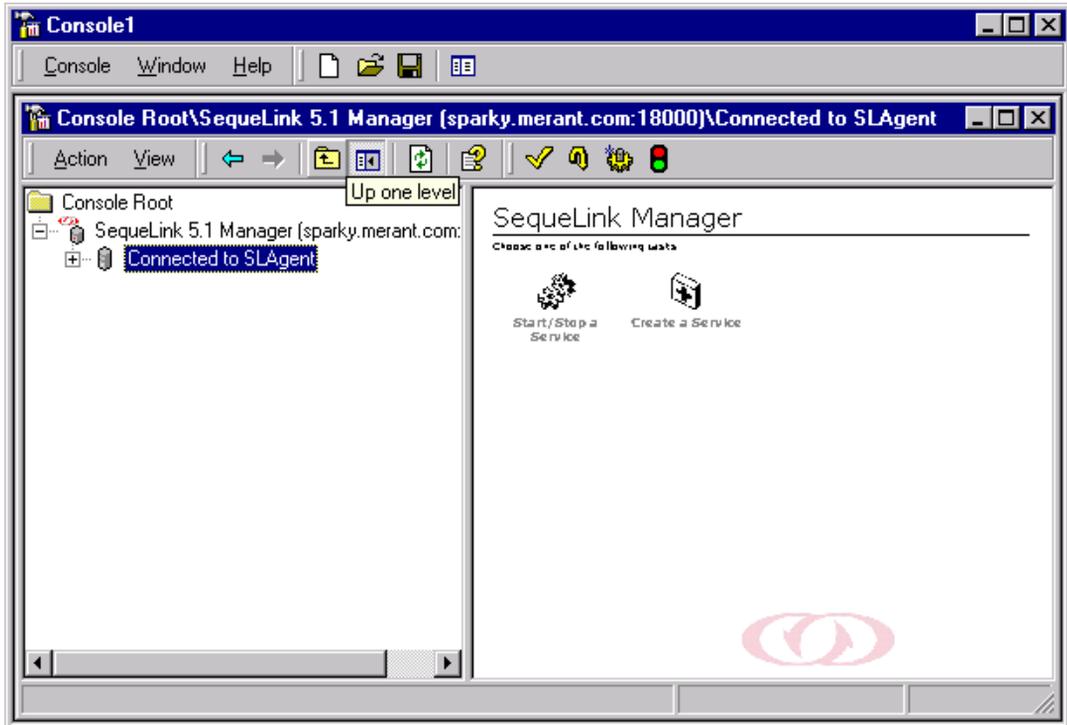
Creating a Server Data Source

This example shows how to use the SequeLink Manager MMC Snap-in to create a server data source.

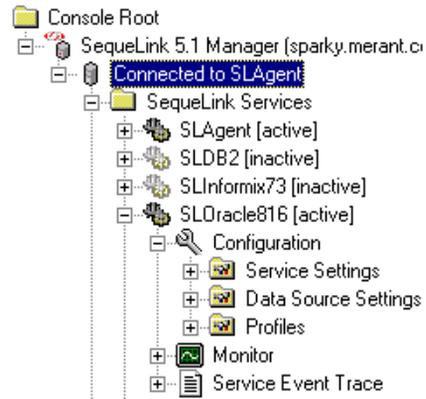
- 1 Start the SequeLink Manager MMC Snap-in by clicking **Start / Run**. In the Run field, type **mmc** and click **OK**. An MMC console window appears.



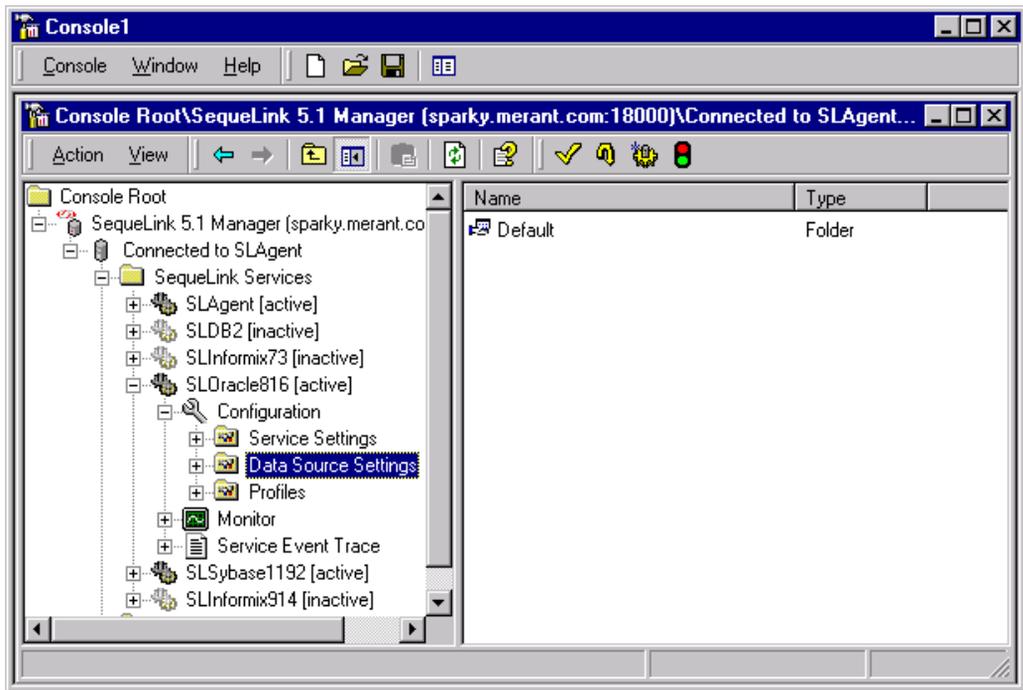
- 2 From the Console menu, select **Console / Open**. Select the SequeLink Manager console you want to start; then, click **Open**. The SequeLink Manager appears in the MMC console window.



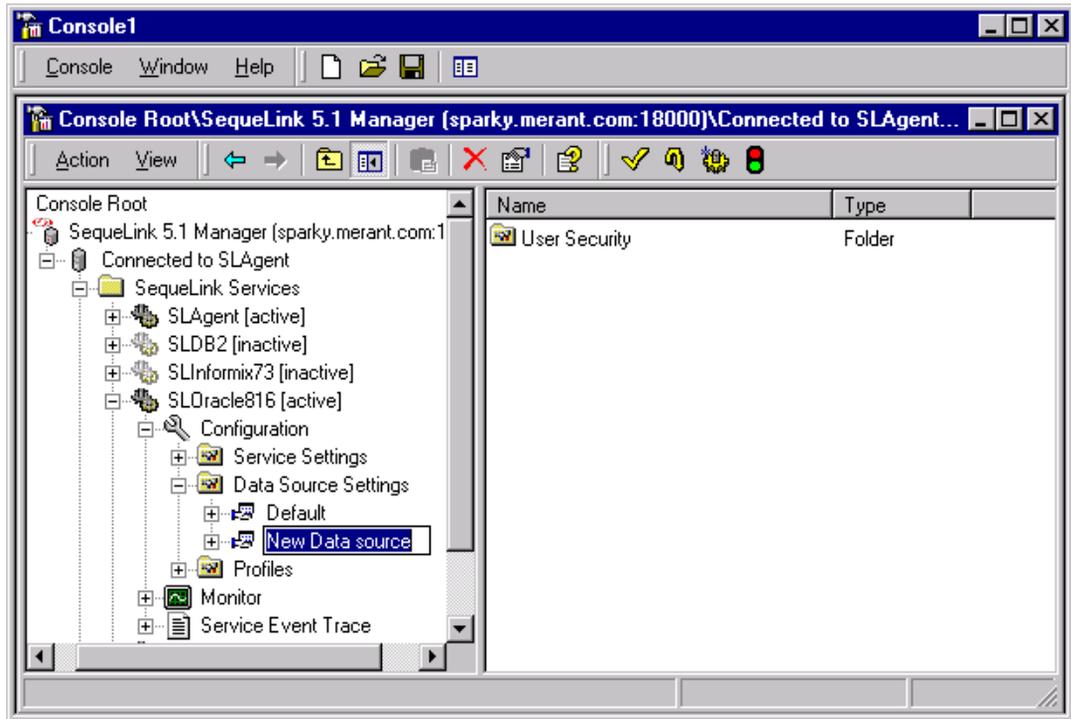
- 3 Open the following nodes:
- Connected to SAgent
 - SequeLink Services
 - SLOracle816
 - Configuration



Then, select the Datasource Settings node to display the existing data sources in the Details pane (right pane).

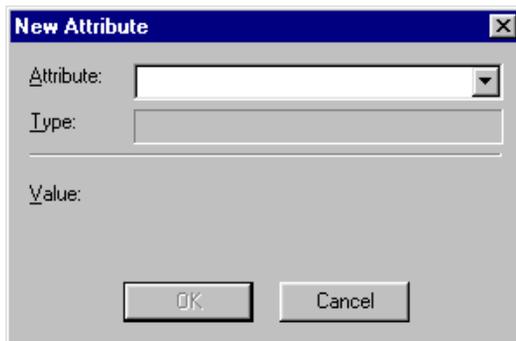


- 4 Right-click the Datasource Settings node and select **New / Datasource**. A new server data source appears in the left pane as an editable field.



- 5 To name the server data source, type `Or8Update` and press ENTER.

- 6 Add the DataSourceApplId attribute to the Or8Update data source. Right-click the **Or8Update** data source and select **New / Attribute**. The New Attribute window appears.



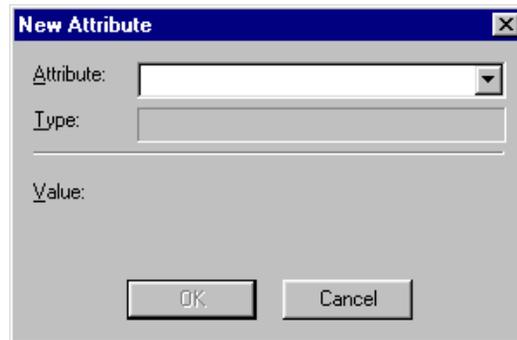
- 7 In the Attribute drop-down list, select **DataSourceApplID**.
- 8 In the Value field, type the application ID the order-entry application will use to identify itself to the SequeLink service. Refer to the *SequeLink Developer's Reference* for information about specifying application IDs for ODBC client applications.
- 9 Click **OK**. The attribute is added to the server data source.

For a description of SequeLink service attributes, refer to the *SequeLink Administrators Guide*.

Modifying the Default Server Data Source

NOTE: This procedure assumes that the SequeLink Manager is running and that you have completed the previous procedure.

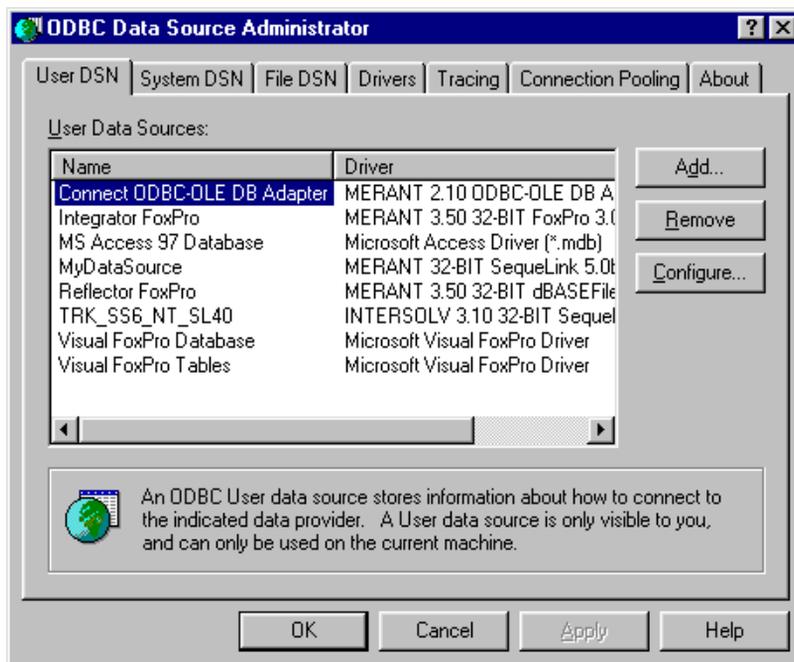
- 1 Select the Datasource Settings node to display the existing data sources in the Details pane (right pane).
- 2 Add the DataSourceReadOnly attribute to the Default data source. Right-click **Default** in the Details pane and select **New / Attribute**. The New Attribute window appears.



- 3 In the Attribute drop-down list, select **DataSourceReadOnly**.
- 4 In the Value field, select **Select**, which means only Select statements are allowed.
- 5 Click **OK**. The attribute is added to the default server data source.

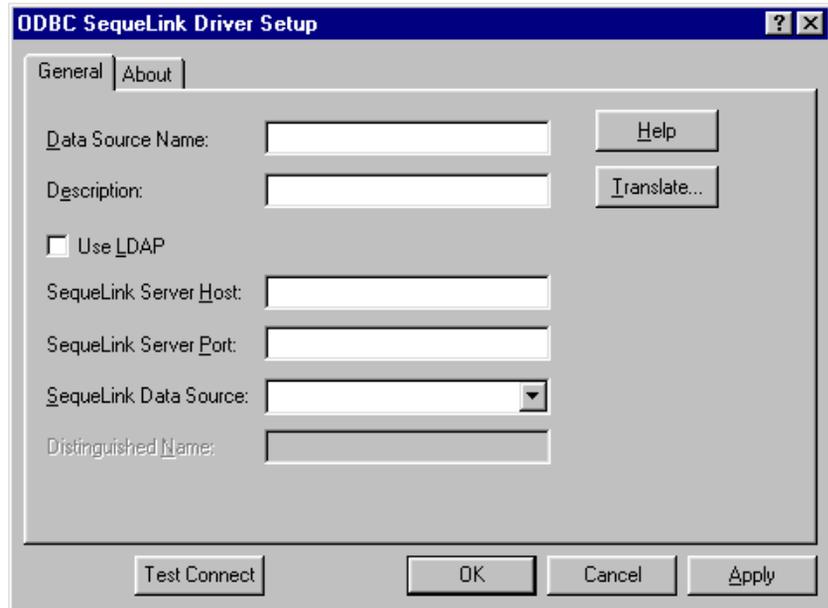
Configuring ODBC Client Data Sources

- 1 Start the ODBC Administrator by clicking **Start**, and then **Programs**. From the Programs menu, select **SequeLink ODBC Client 5.1**, and then select the **ODBC Administrator** application. Click the **User DSN** tab to view a list of existing user data sources.



- 2 Create and configure the first data source by clicking the **Add** button. A list of installed drivers appears. Select **MERANT 32-BIT SequeLink 5.1**; then, click **Finish**.

The ODBC SequeLink Driver Setup window appears.



- 3 Provide the following information; then, click **OK**.

Data Source Name: Type `Or8ReadOnly`.

Description: Type `Oracle 8 read-only connection`. This field is optional.

SequeLink Server Host: Type `OracleDBServer`. This field identifies the TCP/IP host name of the SequeLink service to which you want the SequeLink ODBC Client to connect.

SequeLink Server Port: Type `1996`, which is the default TCP/IP port the SequeLink service is listening on for incoming connection requests. The port that you specify must be the same as the one that was specified for the SequeLink service when the SequeLink Server was installed.

Server Data Source: Leave this field blank. This client data source is to use the default server data source; therefore, a server data source does not have to be specified.

- 4 Create and configure the second data source by clicking the **Add** button. A list of installed drivers appears. Select **MERANT 32-BIT SequeLink 5.1**; then, click **Finish**.

The ODBC SequeLink Driver Setup window appears.

- 5 Provide the following information; then, click **OK**.

Data Source Name: Type `OR8OrderEntry`.

Description: Type `Oracle 8 update connection`. This field is optional.

SequeLink Server Host: Type `OracleDBServer`. This field identifies the TCP/IP host name of the SequeLink service to which you want the SequeLink ODBC Client to connect.

SequeLink Server Port: Type 19996, which is the default TCP/IP port the SequeLink service is listening on for incoming connection requests. The port you specify must be the same as the one that was specified for the SequeLink service when the SequeLink Server was installed; the default is 19996.

Server Data Source: Type OR8Update, which is the name of the server data source that was created for update capabilities and the use of application IDs, or select the value from the drop-down list.

The configuration of SequeLink client and server data sources is complete. The reporting applications can now use the OR8ReadOnly client data source, and the order-entry application can use the OR8OrderEntry client data source.

For information about how to configure the order-entry application to specify an application ID, refer to the *SequeLink Developer's Reference*.

4 Migrating to SequeLink 5.x

SequeLink version 5.x is not compatible with SequeLink 4.5 configurations. SequeLink 4.5 Clients cannot connect to SequeLink 5.x Servers and vice versa. However, these two versions of SequeLink (Client and Server) can be installed side-by-side without conflicting with the operation of one another.

You may want to install both versions of SequeLink side-by-side so that you can implement the migration to SequeLink 5.x without any interruption in your everyday use of SequeLink 4.5. Then, when you have completed the configuration and testing of SequeLink 5.x, you can delete SequeLink 4.5 from your clients and servers.

IMPORTANT: When installing SequeLink 5.x on servers that also have SequeLink 4.5 installed, you must use a TCP/IP port different from the one used for SequeLink 4.5.

4.5 Data Sources Vs. 5.x Data Sources

SequeLink 4.5 has two types of data sources that reside on the client, a SequeLink CAT data source and a data access data source (either ODBC or OLE DB). When configuring SequeLink 4.5, you have to define a SequeLink CAT data source and an ODBC or OLE DB data source; the ODBC or OLE DB data source references the SequeLink CAT data source.

In contrast, SequeLink 5.x has only one data source that resides on the client, which is called a SequeLink Client data source (either ODBC, ADO, or JDBC) and is similar to the SequeLink 4.5 ODBC or OLE DB data source. In addition, SequeLink 5.x has a

server data source. When configuring SequeLink 5.x, a client data source is required for ODBC and ADO configurations; server data sources are optional. **SequeLink CAT data sources no longer exist in SequeLink 5.x.**

The information that was defined in a SequeLink 4.5 data source is either no longer applicable, has been moved to the server data source, or has been moved to the client data source in SequeLink 5.x.

For current, detailed information about migrating, refer to *Migrating to SequeLink 5.0*. You can locate the document by going to:

<http://198.77.127.204/cgi-bin/webcgi.exe?new,kb=DataDirect,TemplateSet=MERANT8>

and searching on the keyword "migrating."

Glossary

ActiveX Data Objects (ADO)	A high level object-oriented database API built on OLE DB.
Client data sources	Minimal ODBC, ADO, or JDBC data sources that store connection instructions to a SequeLink Server.
Database Management System (DBMS)	A layer of software between the physical database and the user. The DBMS manages all access to the database.
data store	The storage device for data a user accesses, such as the data in a database or a file. A data store owns data and exposes its data in a tabular form as a rowset over a native data format. Data stores can include a full SQL DBMS, an ISAM file, or a text file or data stream.
Distinguished Name (DN)	A name that identifies an LDAP entry in an LDAP directory. See also Lightweight Directory Access Protocol (LDAP) .
JDBC	A data access API standard (Java Database Connectivity) for Java-enabled applets, applications, or Web browsers.
Lightweight Directory Access Protocol (LDAP)	A standard protocol for accessing and updating common directory information.
Microsoft Management Console (MMC)	The MMC is a common console framework for management applications.
middleware	Software that mediates the communication between an application and a data store. The middleware provides an interface that manages the differences in the application's and the data store's data formats.
ODBC	Microsoft's Open Database Connectivity (ODBC) specification. The ODBC specification for an Application Programming Interface (API) allows applications to access multiple database systems using Structured Query Language (SQL). For detailed information on ODBC, refer to the Microsoft programming documentation on ODBC.

SequeLink administrator	Typically, the person who configures, manages, and monitors the SequeLink environment.
SequeLink ADO Provider	The data provider that is installed within the SequeLink ADO Client to provide ADO/OLE DB access.
SequeLink Agent	The SequeLink service that acts as a proxy to carry out configuration, management, and monitoring requests from the SequeLink Manager.
SequeLink Client	A SequeLink software component that can be installed on a client machine or on a Web/Application Server. SequeLink ODBC Client provides ODBC access; SequeLink Java Client provides JDBC access; SequeLink ADO Client provides ADO/OLE DB access.
SequeLink configuration file	A file that contains configuration information for SequeLink services.
SequeLink data sources	Optional data sources stored on the SequeLink server containing service-specific instructions that affect how data is accessed by a connection. Centralizing this information on the server, instead of distributing it among hundreds of clients, provides easier management of your entire data access infrastructure.
SequeLink ADO Driver	The driver that is installed with the SequeLink ADO Client to provide ADO access.
SequeLink JDBC Driver	The driver that is installed within the SequeLink Java Client to provide JDBC access.
SequeLink ODBC Driver	The driver that is installed within the SequeLink ODBC Client to provide ODBC access.
SequeLink Manager	A tool that you can use to configure, manage, and monitor your SequeLink environment. The SequeLink Manager is provided as an MMC snap-in on Windows, a command-line tool on Windows and UNIX, and an ISPF panel tool on OS/390.
SequeLink Manager Command-Line Tool	A command-line tool supported on Windows and UNIX platforms that allows you to configure, manage, and monitor your SequeLink environment.
SequeLink Manager for OS/390	An ISPF panel tool supported on OS/390 that allows you to configure, manage, and monitor your SequeLink Server for OS/390 services locally from an OS/390 machine.

SequeLink Manager Snap-in	A GUI-based tool supported on Windows platforms that allows you to configure, manage, and monitor your SequeLink environment. Before you can use the SequeLink Manager Snap-in remotely, you must add it to the MMC.
SequeLink profile	Predefined profiles that you can use for monitoring your SequeLink environment, such as viewing details about active services, viewing active sessions, and requesting information about traced events.
SequeLink Server	The SequeLink software component that is installed on the server to provide data access services from client applications to data stores.
SequeLink service	SequeLink provides the following service types: <ul style="list-style-type: none">■ <i>SequeLink data access services</i> handle data access requests from any SequeLink client. Multiple SequeLink data access services can run on the same SequeLink server. For example, SequeLink Server for Oracle and SequeLink Server for Microsoft SQL Server can run side-by-side on the same machine.■ <i>SequeLink Agent services</i> act as a proxy to carry out configuration, management, and monitoring requests from any SequeLink Manager. The SequeLink Agent can service multiple SequeLink services on the same SequeLink server.
SequeLink service template	A template that contains predefined attributes for a particular service type, for example, Oracle8.
SQL	Structured Query Language. A language used by relational databases to query, update, and manage data.

Index

A

- ADO data sources 89
- algorithms
 - asymmetric cryptographic 66
 - message digest 67
 - symmetric cryptographic 67
- application IDs, about 64
- architecture
 - n-tier 33
 - two-tier 31
- asymmetric cryptographic algorithms 66
- authentication 60
- authorization 61

B

- books, ordering printed 12

C

- CAT data sources 89
- cipher suites 67
- client data sources 30
 - about 30
 - configuring with Quick Install images 47
- configuration
 - information you need before you start 42
 - planning for 41
 - requirements 42

- configuring
 - service monitoring 72
 - session monitoring 73
 - statement monitoring 73
- connection model
 - default behavior
 - OS/390 53
 - Windows and UNIX 49
 - planning 58
- connections, configuring for SequeLink Server 44
- contacting Technical Support 14
- conventions used in this book 9

D

- data access services
 - about 22
 - planning 47
- data privacy
 - about 65
 - OS/390 54
 - Windows and UNIX 50
- data scrambling 65
- data sources
 - configuring client configurations 47
 - default behavior 50, 54
 - differences from earlier versions 89
 - Quick Install images 47
- data stores, linking to 29
- database logon 64
- DB2 default behavior
 - OS/390 55
 - Windows and UNIX 51

- debug default behavior
 - OS/390 53
 - Windows and UNIX 50
- default behavior
 - connection model
 - OS/390 53
 - Windows and UNIX 49
 - data privacy
 - OS/390 54
 - Windows and UNIX 50
 - DB2
 - OS/390 55
 - Windows and UNIX 51
 - debug
 - OS/390 53
 - Windows and UNIX 50
 - event handling
 - OS/390 53
 - Windows and UNIX 49
 - examples of changing 56
 - general data source
 - OS/390 54
 - Windows and UNIX 50
 - general service
 - OS/390 54
 - Windows and UNIX 50
 - Informix 51
 - Microsoft SQL Server 52
 - Oracle 52
 - security
 - OS/390 53
 - Windows and UNIX 50
 - Sybase 52
 - distributed transactions 17
 - documentation
 - order form 13
 - ordering hard-copy books 12

E

- encryption 65
- event handling default behavior
 - OS/390 53
 - Windows and UNIX 49
- event trace file
 - OS/390 53
 - Windows and UNIX 49
- event tracing, planning 72

G

- general service default behavior
 - OS/390 54
 - Windows and UNIX 50
- glossary 91

I

- IIOp (Internet Inter-ORB-Protocol) 18
- Informix default behavior 51
- interoperability 17

J

- JDBC 2.0 Optional Package 17
- JDBC data sources 89

L

LDAP

- configuring connections for SequeLink Server 44
- overview of use 17
- linking to data stores 29

M

- message digest algorithms 67
- migrating from earlier versions 89
- monitoring activities
 - by service 72
 - by session within a service 73
 - by statement within a service 73
- monitoring, planning 72

N

- NT Integrated security 18

O

- ODBC data sources 89
- OLE DB data sources 89
- Oracle default behavior 52
- ordering hard-copy books 12
- ordering printed books 12

P

- planning 47
 - client data sources 43
 - configuration 41

- data access services 47
- information you need before you start 42
- monitoring and event tracing 72
- scenarios 35
- SequeLink Client configurations 47
- SequeLink connection model 58
- SequeLink data access services 47
- system administration 70

printed books

- order form 13
- ordering 12

privacy, data

- OS/390 54
- Windows and UNIX 50

Q

- Quick Install images
 - configuring 46
 - including data sources in 47

R

- ReadOnly 65

S

- scalability 18
- scenarios, planning 35
- Secure Socket Layer (SSL) 66
- security
 - application IDs 64
 - authentication 60
 - authorization 61
 - data scrambling 65
 - database logon 64
 - default behavior

- OS/390 53
 - Windows and UNIX 50
- infiltrations 66
- NT Integrated 18
- overview 18
- planning
 - OS/390 68
 - Windows and UNIX 67
- RACF 18
- ReadOnly 65
- SequeLink
 - n-tier architecture 33
 - two-tier architecture 31
- SequeLink ADO Client, planning
 - configuration 46
- SequeLink Agent, about 22
- SequeLink Client
 - configuration requirements 42
 - creating a common installation image 46
- SequeLink Java Client
 - changes for data sources 89
 - encryption 65
 - JDBC data sources 45
- SequeLink ODBC Client
 - changes for data sources 89
 - example of changing default behavior 57
 - planning configuration 46
- SequeLink Server
 - authentication 60
 - authorization 61
 - configuration requirements 42
 - configuring connections 44
- SequeLink services 24
- server data sources 30
- Service monitoring 72
- ServiceEncryptionAlgorithm 65
- Session monitoring 73
- side-by-side installation 89
- Statement monitoring 73
- SupportNet 14
- Sybase default behavior 52
- symmetric cryptographic algorithms 67
- system administration 70

T

- Technical Support, contacting 14
- transactions, distributed 17